

Expert Reflection

Cautionary Notes on Disinformation and the Origins of Distrust

By [Yochai Benkler](#) · October 22, 2019

Table of Contents

- [Anxiety, in perspective](#)
- [Evidence of activity vs. evidence of effect](#)
- [Origins of distrust and disaffection: Putting disinformation in context of more structural drivers](#)

Benkler, Yochai. 2019. Cautionary Notes on Disinformation and the Origins of Distrust. Social Science Research Council, MediaWell. <https://mediawell.ssrc.org/expert-reflections/cautionary-notes-on-disinformation-benkler/>. DOI: 10.35650/MD.2004.d.2019.

Anxiety, in perspective

Something is profoundly wrong with American democracy when a lead opinion writer, in one of the country's flagship newspapers, calls the Senate majority leader "[a Russian asset](#)." The past three years have seen an explosion of academic, journalistic, and government research and analysis of disinformation, [propaganda](#), and radicalizing hate speech. As we enter the 2020 election season, concerns over foreign election interference, social bots, and other forms of disinformation are flooding our media system. In part, they reflect the significant findings of the range of scientific, journalistic, and governmental investigations that have uncovered new pathways of propaganda and disinformation used by a range of actors—from domestic politicians, through commercial clickbait purveyors, to foreign agents. But they also reflect a visceral response to political outcomes that are far afield of what would have been considered the normal range of political contestation a decade ago. The shock of these dramatic political outcomes—the success of the Leave campaign in the UK, the electoral success of Donald Trump—has, I believe, undergirded the [narratives](#) that these are not the result of a legitimate

underlying transformation in popular sentiment, but are instead the result of overwhelming technological change and sinister foreign intervention.

I offer two cautionary notes to put our anxiety in perspective. At root, we must continue our work, but be more circumspect about whether what we are finding is the proximate or root cause of the crisis of democratic societies in the early twenty-first century. There is important work to be done, both qualitative and quantitative, that must continue to seek to identify new practices that have the potential to undermine democratic elections and public discourse. Identifying Russian operations, for example, is plainly important, technically difficult, and discloses activities that we are interested in as a public. Nonetheless, evidence of action is not evidence of impact—and in the case of Russian interference in the 2016 US presidential election, there is little evidence that these efforts actually made a meaningful difference. As we enter a new election cycle—as reports of Iranian interventions join continued concern with Russia—it is critical that we not overstate the success of these campaigns. Nonstop coverage of propaganda efforts and speculation about their impact, without actual evidence to support that impact, feeds the loss of trust in our institutions to a greater extent than the facts warrant. No specific electoral outcome better serves Russia's standard propaganda agenda—disorientation and loss of trust in institutions generally—than the generalization of disorientation and delegitimation of institutions represented by the present pitch of the debate over election interference.

At the same time, there is enough evidence to believe that information operations are underway, the risk is real, and the only way we will actually know whether our democracy has been compromised is to pass legislation that forces the major platforms, in particular Facebook, to make their data effectively available for independent investigation of their practices and their real impact. It has been two and a half years since the 2016 presidential election, a year and a half since Facebook announced its cooperation with the Social Science Research Council and Social Science One on issues of research data sharing, and four months since Facebook launched its Ads database. Yet technical difficulties and concerns over privacy protections keep the information the company shares extremely limited and nearly impossible to work with in ways that actually allow causal analysis about which interventions—whether ads or astroturf campaigns—actually had an impact on the election. And Facebook is further along in sharing its information than, say, Google. I fear that nothing short of legally enforced data access for independent researchers, coupled with legally enforceable privacy protection for users, will get us the kind of information we need to be confident that our election integrity has not been compromised.

In the second part of this piece, I turn to a more basic version of the same question, a version that cannot be answered by any policy directed at Facebook or any other technology firm: To what extent can we claim that disinformation and propaganda are a core cause of the crisis of democratic society? I suggest that we may need to understand the political confusion less as a function of intentional, discrete propagandist interventions (though it is those too), and more as one facet of the more basic dynamic of the past decade: the collapse of neoliberalism in the Great Recession and the struggle to define what the new settlement that replaces it must become. I suggest that while individual entrepreneurs took advantage of the dislocation and anxiety caused by these conditions—be they media personalities like Rush Limbaugh and or Rupert Murdoch or political opportunists like Donald Trump—the driving force of

the epistemic crisis is that elite institutions, including mainstream media, in fact failed the majority of the people. Throughout the neoliberal period, elite consensus implemented policies and propagated narratives that underwrote and legitimated the rise of a small oligarchic elite at the expense of delivering economic insecurity to the many. These material drivers of justified distrust were compounded by profound changes in political culture that drove movements, on both the right and the left, to reject the authority structures of mid-twentieth-century high modernism.

Evidence of activity vs. evidence of effect

Did the Russians hack the 2016 US presidential election, and will they hack the 2020 election? The answers depend on what we mean by that question. One interpretation is: “Did actors operating for the Russian state conduct information operations aimed at the United States and intended to exacerbate preexisting divisions in general and help Donald Trump get elected, and will they do so again?” Another interpretation of the same question is: “Did Russian activity make a measurable contribution to Donald Trump’s victory?” The former calls for evidence of activity that is reasonably tied to Russia. This kind of evidence can be very technical and difficult to obtain, but once practices are discovered and tied to Russia, the mission is accomplished. The latter calls for much more detailed evidence and more complex analysis, because teasing out the relative contribution of observed activity to as complex a social fact as the election of Donald Trump is a much steeper hill to climb. Having struggled mightily with both aspects of the question in my own work with colleagues, it seems quite clear to me that the answer to “Did Russia target the US?” is yes. By contrast, evidence that Russian operations made a meaningful difference in an admittedly close election is thin and speculative. And what is true of Russian propaganda is equally true of targeted advertising on Facebook generally, and in particular the claim that it was intensive targeted advertising on Facebook, aimed to suppress Clinton voters and turn out Trump voters, that gave Trump the winning edge.

The facts of Russian operations against the US 2016 election are established by four major lines of inquiry. The first is governmental—the Mueller indictments against various Russian entities and individuals, the House and Senate investigations, and in particular the congressional testimony of Facebook, Twitter, and Google identifying Russian accounts that purchased advertising as well as Russian-operated accounts masquerading as American accounts aimed across the political divide. The second is a steady flow of investigative journalism by diverse sources, from the Daily Beast to the *Wall Street Journal*. The third are distinct reports by cybersecurity firms on specific claims—in particular, the Russian origins of the Democratic National Committee (DNC) emails hack. And the fourth are academic reports that look at specific campaigns and trace their Russian origins. It is not impossible that all these diverse independent sources, doing work over a period of three years, are simply following each other’s lead and reflecting a society-wide groupthink. But it is much less likely that all these diverse avenues and institutional settings have produced the same error than that Russia in fact ran information operations aimed at the American political system.

That's a far cry from evidence that Russian information operations in fact made a difference. While Hillary Clinton's candidacy was dominated by coverage of "emails," little of that coverage was focused on emails that resulted from the two Russian hacks—the DNC and Podesta emails. Rather, most of the coverage involved the private server, State Department releases, and the James Comey announcements, most prominently the announcement in the week before the election that the FBI investigation was being reopened because of a computer found in a separate investigation of Anthony Weiner, and then the announcement that the investigation was again closed because there was nothing of note on the computer. Coverage of these latter announcements dominated all other topics in the days before the election. Looking at the Facebook and Twitter campaigns alleged in the Mueller indictment of the Russian Internet Research Agency, the primary actor in the online sockpuppet (human-operated fake accounts) campaigns, my colleagues and I found that the alleged interventions generally followed an uptick in interest in a topic—such as voter fraud—where the origin of the interest was something that Trump had said and was widely covered in major media, both mainstream and right-wing oriented, like Fox News or Breitbart. Russian sockpuppet posts came a few days after the peak of interest. Even where we did find traces of Russian origins in campaigns that did make it into the mainstream, the propaganda pipeline ran through Infowars, Drudge, and Fox News. That is, the primary vectors of influence were willing amplification of the information operations by the mainstays of the domestic American outrage industry. Nor have any of the agencies reporting Russian efforts to gain access to state-based electoral institutions made any claim that any of the actual voting procedures, much less machines, were compromised. There is one suggestive recent study that finds a correlation between the Twitter activity of known Russian accounts and changes in polling data, where a rise in activity is a leading indicator of a rise in polls, but the authors appropriately acknowledge that this mass-scale observation cannot justify individual causal claims (e.g., that exposure led to a change in views) or account for broader media ecosystem effects.

What is true of Russian interference specifically and foreign interference more generally is also true of Facebook targeted advertising. That risk was embodied by Cambridge Analytica—the now defunct and disgraced company that at one point claimed credit for both Brexit and Trump's victory. The idea was that individually targeted campaigns—both paid advertisements and astroturf sockpuppets—successfully influenced people's behavior and beliefs, operating most directly through turnout interventions. There is good reason to think that Cambridge Analytica's specific intervention—psychographically informed targeted advertising—was not used, and more importantly, that its proven impact is so small that it is impossible to credit with shifting even such a tightly contested election as the 2016 US presidential election. But we lack the data to know whether the much larger investments of the political campaigns in Facebook advertising in fact made a meaningful difference, and if so, whether the difference was in get-out-the-vote efforts to their own supporters or in disinformation campaigns aimed to suppress the major pro-Clinton voter blocks—women, young, and Black voters.

It is this kind of evidence of causation and effect, rather than purely of activity, that we now need urgently. That we know we are the target of foreign information operations is important, and justifies investing in identifying and tracing these kinds of moves. That we know that Facebook sells individually targeted advertising that could be used to narrowcast disinformation to select groups in ways that may

be illegal is also important. But neither of these beliefs justifies loss of trust in our electoral system and its outcomes unless we know that these vectors of disinformation, rather than some other factor, are changing political outcomes. Because if at the end of the day all social science and formal investigation can do is raise anxieties and put us in a position to have better evidence on which to base our fears, but no ability to distinguish between fears that are actually justified by impact and peripheral phenomena, we will end up embracing and supporting policies that are unnecessarily restrictive and problematic.

The studies closest to providing the kind of data we need were published early in 2019 in *Science*. Both were focused on “fake news” rather than on Russian propaganda or targeted advertising, but their methods point to the kind of individualized and rich demographic data we would need to assess impact. In one, [Grinberg et al.](#) reported on tweeting patterns of over 16,000 Twitter accounts that they were able to link to real-world registered voters. What they found was that 0.1 percent of the users were responsible for sharing 80 percent of the fake news content, while 1 percent consumed 80 percent of the fake news content. Americans over 65, and in particular those who read right-wing sites, were especially likely to consume fake news. Most users were exposed to very little fake news on Twitter. The second paper, by [Guess et al.](#), looked at the Facebook usage patterns of a representative sample of 3,500 survey respondents and came to a highly congruent observation. Just as Grinberg et al. found for Twitter, Guess et al. found that sharing of fake news on Facebook was rare, that it was concentrated among conservative and very conservative users, and that it was strongly associated with age, in particular among users older than 65. Though neither study drew this conclusion, given how rare fake news exposure and sharing was, and given its concentration among older Americans and politically engaged conservatives, the likelihood that fake news moved these already politically mobilized users with clear political preferences is small. And we can only make that assessment by looking at population-level or representative samples of users and their patterns, rather than looking for evidence that the bad information is being produced out there. A study that sets out to find fake news purveyors whose stories were widely shared will find them. But such a study would not shed light on how important fake news was to the election in the way that these studies could. It is this kind of user-side, detailed information from the whole population or representative samples that we need in order to understand how important and threatening foreign intervention or targeted advertising are to democratic elections.

The platform companies, particularly Facebook and Google, have all the relevant data and much more than has been available even to the best-designed of studies I noted. Facebook, for example, has a full record of every exposure of every user in the run-up to the 2016 election. In principle, it should be possible to look at every single instance of exposure to the different sources of illegal influence—Russian campaigns, for example, or sustained campaigns that might fall under election law but are not declared as such. And in principle, it should be feasible to build likely voter models to assess whether exposure to these kinds of illegal or questionable information operations has an impact on individual-level behavior, how large an impact, and in which states. Google has a full record of exposure to YouTube videos and advertisements. And like Facebook, it possesses a wealth of individualized information about users that could be used to make a similar assessment. There exists, in other words, data out there that could answer quite conclusively whether Russian propaganda, “fake news purveyors,” or campaign-driven targeted advertising actually made a difference in the outcomes of the major political upsets of 2016. But

no independent investigators have access to that data.

As one conceives of what such a study might look like, the privacy concerns are genuine, and the difficulties of negotiating private arrangements are significant. But we must also recognize that the firms have strong incentives to remain opaque about their role in the elections. If careful study shows that large investment in Facebook advertising and manipulation made the marginal difference in, say, Donald Trump's election, the company risks the ire of its Democratic users. If, by contrast, a careful study shows that campaigns on the platform had no meaningful impact, then the company's core business model—selling targeted ads on the promise that they will sway voters and consumers—will be shown to be ineffective. So genuine privacy concerns align with a deep conflict of interest between the firms' interests and what the public needs. Perhaps over time the Social Science One initiative will be able to deliver data at a sufficient level of granularity to make clear causal inferences. Even then, the dataset would cover only one important pathway. What is likely to be necessary to solve the problem at a systemic level is a public process, with enforceable sanctions against firms that do not cooperate and against independent researchers who gain access under this process but then breach privacy. Until then, political advertising, propaganda, and disinformation campaigns will continue to be a faith-based process, and our efforts to combat them too will be based on reasonable anxiety about the possibility of impact rather than reasoned response about a known effect.

Origins of distrust and disaffection: Putting disinformation in context of more structural drivers

Mary Shelley's *Frankenstein* was published a couple of years after the last lacemaking machine was shattered by the Luddites. Anxiety over technology has, it seems, been a lifelong companion of modernity. In this sense, the focus on Macedonian teenagers manipulating the Facebook advertising algorithm, Russian hackers, bots, and sockpuppets, and Cambridge Analytica-style behavioral manipulation in the wake of Brexit and Trump's election is unsurprising. But the politically inflected, asymmetric patterns of disinformation and propaganda in the American case should make us skeptical of these usual suspects. Instead, we need to interpret what we see, as well as the success of any given propagandist, in the context of long-term patterns of loss of trust in institutions, including mainstream media, and the deep alienation of the past decade since the Great Recession. Here, I sketch out these longer-term patterns and outline their structural drivers. Underscoring these doesn't mean that one or another source of disinformation is unimportant. But it does mean that solutions that focus purely on fighting symptoms and opportunists taking advantage of the disruption and disorientation of the past decade will miss their mark in the long term.

Our present experience of epistemic crisis cannot be separated from the much broader and deeper trend of loss of trust in institutions generally. When we look at the longest time series regarding trust in any institution—Pew Research Center's series on trust in the federal government—we see that most of the decline in trust occurred between 1964 and 1980. Pew's long series shows that this change was not an

intergenerational shift. There is no meaningful difference between the sharp drop in trust among the “greatest,” “silent,” and “boomer” generations. And the drop from 77 percent who trust government in 1964 to 28 percent in 1980 dwarfs the remaining, up-and-down noisy drop from 28 percent in 1980 to 19 percent from 2014 through 2017. Gallup’s long-term data series, starting from 1973 (already well down the decline curve if we use the longer-term trust-in-government series), suggests an across-the-board decline in which trust in media does not stand out. Only the military and small business fared well over the period from 1973 to the present. Big business and banks, labor unions, public schools and the health-care system, the presidency and Congress, the criminal justice system, organized religion—all lost trust significantly, and most no less or more than newspapers. Moreover, as the Organisation for Economic Co-operation and Development documented, loss of trust in government is widespread in contemporary democracies. The secular decline of diverse institutions, the sharp change in the 1960s and 1970s, and the parallels in other countries require caution before proceeding on the assumption that discrete, intentional actors operating specifically on one institutional system (be they domestic supporters of the outrage industry like Rupert Murdoch and Robert Mercer or foreign intelligence services) were the core driving factor. Broad structural changes could well be accompanied, perhaps accelerated, by distinct individual and corporate actions intended to capitalize on them, and these will appear in retrospect to have made all the difference.

What happened in the 1960s and 1970s that could have caused this nearly across-the-board decline in trust in institutions? One part of the answer focuses on material origins: the end of the “Golden Age of Capitalism,” its replacement by neoliberalism, and the spread of broad-based economic insecurity that followed. The other part of the answer is anchored more directly in political culture, and encompasses rejection of authority and elite power structures by both left-wing elements—the antiwar movement, the New Left, and the women’s movement—and right-wing elements, in the United States particularly white identity voters and the newly politicized Christian fundamentalists.

The “Golden Age of Capitalism” that began after World War II was a unique, large-scale global event typified by high growth rates across the industrialized world from recovery investment, at a time when war-derived solidarism pushed institutions, even in the United States, that led to broad-based economic security and declining inequality. Whether one locates the end of the “Golden Age” in the profit squeeze of the late 1960s, the collapse of Bretton Woods in 1971, or the oil shocks of 1973 and 1979, it is clear that by the 1970s widespread economic anxiety of a form utterly unrecognizable in the early 1960s had settled in.

The Great Inflation of the 1970s marked the end of broad public acceptance of the central role of government stewardship of the economy, the Keynesianism and dirigisme that marked the era from the end of World War II to the early 1970s. That loss of trust in government stewardship, driven by the economic disruption, created an opening for neoliberalism to emerge as the dominant frame for organizing the economy. Rooted in the work of Hayek, Friedman, and others in the Mont Pelerin Society, and implemented politically with the victories of Margaret Thatcher in 1979 and Ronald Reagan in 1980, neoliberalism, or what for a while was called the Washington Consensus, became the standard answer among professional elites for how to manage the relationship between the state and the market.

Privatization, deregulation, lower taxes, and the free movement of goods, finance, and people marked a dramatic reversal of the policies of the “Golden Age of Capitalism,” shrinking the role of the state significantly and embracing a near-religious belief in the wisdom and beneficence of markets.

In retrospect, we now know that the promises of that program—growth and trickle-down prosperity for all—failed to materialize. Instead, embrace of the neoliberal program resulted in a dramatic rise in inequality and growing economic insecurity for the broad middle classes in the advanced economies, nowhere more so than in the United States. Median income, which had grown in line with productivity growth from 1946 to 1973, flatlined, and in the US, median real wages have effectively not grown since 1973. Meanwhile, a broad range of policy changes privatized risk, shifting its burden onto a middle class unable to shoulder it. Diseases of despair—alcoholism, opioid addiction, and suicide—in turn, have made white Americans with a high school education or less the only population in the advanced world to have seen its life expectancy decline since the early 1990s. Small wonder that this broad middle class finds itself alienated and ready to hear from populists who blame immigrants and elites for their woes. Indeed, research in the past few years across diverse countries following the Great Recession, suggests a strong and positive association between economic insecurity (whether operationalized as exposure to Chinese trade in the US, increases in the rate of unemployment, or risk of automation in Europe) and vote share for anti-establishment populists, particularly of the far-right variety.

The basic point is that under conditions of economic threat and uncertainty, people tend to lose trust in elites of all stripes, since they seem to be leading them astray. Mainstream media, for all its professional focus on fact checking and objectivity, was an integral part of that same elite consensus throughout this period, and therefore was never able to deliver a competent and critical perspective to help people understand why it was that the neoliberal program was causing their misery. Instead, media coverage was dominated by the kind of reporting long decried by communications scholarship: a steady flow of pat distractions—horse-race coverage of elections, celebrity scandals, and patriotic pablum in time of war.

The second part of the answer is more directly rooted in politics and changes in political culture. The sharp inflection point in levels of trust in government in the United States in 1964 suggests that the Civil Rights Acts and the Vietnam War played a crucial role. Here, the civil rights and women’s movements emphasized the repressive and unjust attributes of then-prevailing institutions, leading to significant loss of trust among their members; the New Left rejected most institutions and located the individual and self-actualization at the core of moral concern; the antiwar movement rejected the authority of the national security state; and the consumer and environmental movements attacked crony capitalism for its harmful impact on the human and natural environment. All these combined to drive trust in authority-based institutions down on the left and center left. In the 1970s, this meant that when the neoliberal program of deregulation and privatization was introduced, there was little resistance from center-left parties and no coherent alternative for defining the role of the state in the economy beyond mainline socialism. Indeed, Senator Ted Kennedy and later President Jimmy Carter, egged on by the Nader Raiders and the consumer movement more generally, were at the forefront of the first wave of deregulation—of airlines, trucking, and banking. On the right, the backlash of Southern white identity voters against the civil rights movement was harnessed politically by Nixon’s Southern Strategy. Christian fundamentalists became

politicized in the 1970s, in part in reaction to the women's movement and in part in reaction to the rights revolution that pushed against religion in public life, and were harnessed in Reagan's embrace of the Moral Majority. These created large subsets of the Republican coalition ready to distrust elites and harnessed by the antistate ideology of neoliberalism captured vividly by Ronald Reagan's dreaded bogeyman, evoked by the words, "I'm from the government and I'm here to help."

Complementing these broad changes in political coalitions, the successes of consumer, worker, and environmental campaigns in the 1960s drove the rise of organized business. Beginning in the early 1970s, business organizations dramatically increased their lobbying investments and cooperation, and not only strategically funded politicians to make sure policies followed the demands of business and the wealthy, but also deployed "merchants of doubt" in strategic attacks on political, academic, and media institutions that were not compliant with business interests. Mainstream media portrayals of the world—largely in terms congruent with elite views—widely diverged from the perspectives of critics on both sides of the political map. Declining trust in institutions in each of these distinct segments of the population was, in many cases, a reasonable response to institutions whose actual functioning fell far short of what those segments needed or expected, or that had been corrupted or disrupted as a result of the political process.

Taking the material and political dimensions of the answer together begins to point us toward an answer to the question of why we are experiencing an epistemic crisis now, across many (but not all) democratic or recently democratized countries. I have used the experience of the United States to offer an answer rooted in political economy and culture, rather than technological shock. It is possible that in some other countries, technological shock played a more prominent role. But the consistent correlation between measures of economic insecurity and support for populist parties in Europe, as well as the substantial diversity in the degree to which different populations in these equally technologically enabled democracies in fact have lost trust in institutions and governing elites, suggest that careful analysis will find the political economy and cultural change have played a central role in other countries as well.

The neoliberal program—deregulation, privatization, low taxes, and free movement of goods, capital, and labor—promised economic dynamism in exchange for economic security and enhanced consumer sovereignty and entrepreneurial opportunity in exchange for social solidarity. The reality it delivered was far from that promise, and its failure came to a head with the Great Recession. Nationalist politicians have been harnessing anxieties over economic insecurity to the long-simmering anxieties over ethnic and racial identity, offering charismatic authority to replace technocratic authority (expertise), and atavistic solidarity as a way of pointing the finger at someone else, where the unattainable ideal of self-actualization had framed economic insecurity as the personal failure of those who suffered it. Elites, responsible not only for enacting and legitimating liberalization and globalization on the economic side of governing but also cosmopolitanism and pluralism on the identity side, present a challenge to both parts of the economic nationalism of the new right. As such, elite institutions generally—not only professional media, but also academia, science, the professions, and civil servants and expert agencies—have come to serve as the active oppressive "other" of ethnonationalist politics.

Barring distinct evidence of impact of the kind I outlined in the first half of this essay, it seems likely that the shared, global crisis of the neoliberal settlement since the Great Recession, and not technology and disinformation, is driving what we experience as epistemic crisis. When we study disinformation, therefore, we need to understand our work in the context of these broader destabilization dynamics, rather than imagining that what we are observing is a distinct, technologically created new challenge. More importantly, while I do not object to efforts to tamp down the flames with various propaganda-specific interventions—say, reducing the advertising revenue flow to clickbait fabricators who harness political outrage for quick commercial gains or identifying and interdicting Russian propaganda efforts—it's important to understand that these are, at best, emergency care, and often are no more than palliatives. The crisis of democratic societies may be helped along by disinformation and propaganda, and certainly is fanned and harnessed by political opportunists, but it is fueled by a decades-long extractive and distorted political economy.

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](#).