

*Penultimate draft. Please cite final version forthcoming in Critical Review. A Journal of Politics and Society.*

*Étienne Brown*

## PROPAGANDA, MISINFORMATION, AND THE EPISTEMIC VALUE OF DEMOCRACY

**ABSTRACT:** *If citizens are to make enlightened collective decisions, they need to rely on true factual beliefs, but misinformation impairs their ability to do so. Although some cases of misinformation are deliberate and amount to propaganda, cases of inadvertent misinformation are just as problematic in affecting the beliefs and behavior of democratic citizens. A review of empirical evidence suggests that this is a serious problem that cannot entirely be corrected by means of deliberation.*

**Keywords:** *epistemic democracy; fake news; misinformation; political epistemology; propaganda.*

Étienne Brown, [etienne.brown@philosophy.ox.ac.uk](mailto:etienne.brown@philosophy.ox.ac.uk), 22 St. Albans Villas, London, NW5 1QU, United Kingdom, expresses gratitude for the suggestions of audience members at the Montréal conference on epistemic democracy and institutional design (2017), the Paris workshop on populism (2017), and the London Conference on political epistemology (2018), and to Kristoffer Ahlstrom-Vij, Elizabeth Anderson, Boudewijn de Bruin, Paul Boswell, Jason Brennan, Jeroen de Ridder, Marc-Antoine Dilhac, Jason D’Cruz, Catherine Elgin, Jens Gillissen, Alexander Guerrero, Jeffrey Friedman, Michael Hannon, Richard Healy, Jonathan Jenkins Ichikawa, Hélène Landemore, Neil Levy, Robin McKenna, Pierre-Yves Néron, Angie Pepper, Andrei Poama, Emily Sullivan, and Zoe P. Williams for insightful comments. This research was supported by a Social Sciences and Humanities Research Council of Canada (SSHRC) postdoctoral fellowship held at the Oxford Uehiro Centre for Practical Ethics.

Is democracy fundamentally threatened by those who aim to mislead the public? That has been a perennial worry. In the *Protagoras* (317 a-b), Plato writes that when talented orators try to conceal their art, “the masses, needless to say, perceive nothing, but merely sing the tune their leaders announce.” In the *Politics* (1304 b, 20-25, and 1310 a, 1-5), Aristotle is even more explicit about the dangers of demagoguery: “revolutions in democracies,” he explains, “are generally caused by the intemperance of demagogues,” who “are always cutting the city in two” and lead the masses astray. In a recent contribution to the study of propaganda, Jason Stanley (2015, 27) argues that “a certain form of propaganda, associated with demagogues, poses an existential threat to liberal democracies.” More specifically, he defends the claim that propaganda poses an “obvious threat to the *epistemic conception* of democracy, championed by the philosopher David Estlund and the political scientist Hélène Landemore” (ibid., 12).

Stanley’s version of the argument can be attacked on several fronts. First, it may be that his case is overstated, such that we do not have any serious reason to believe that propaganda is effective. He provides little persuasive empirical evidence on this score. Moreover, even if we believe that under the right conditions, propaganda is effective, it may not pose a significant challenge to democratic life. Social scientists have long argued that citizens are politically

ignorant, and he does not show that a deceived public is more dangerous to the epistemic potential of democracy than an ignorant one. Finally, the pernicious effects of propaganda might be adequately countered by the epistemic value of democratic deliberation, which weeds out “the good arguments, interpretations and information from the bad ones” (Estlund and Landemore 2018, 121).

Against these grounds for skepticism, my suggestion will be that a phenomenon related to propaganda—misinformation—does undermine the epistemic value of democracy. More specifically, I will contend that in real-world mass democracies, misinformation prevents the effective functioning of the mechanisms of collective intelligence conceived of by epistemic democrats, and more specifically by Landemore. If citizens are to make enlightened collective decisions, they need to rely on true factual beliefs, but misinformation impairs their ability to do so. As will become apparent, however, my claim is distinct from Stanley’s, for I contend that all cases of misinformation—not merely those that qualify as propaganda—threaten the epistemic potential of democratic decision making.

My argument will proceed as follows. Section 1 discusses the relationship between propaganda and misinformation and suggest that non-propagandistic misinformation is as problematic as what I call cognitive propaganda. If our concern is with the epistemic value of democracy, then, there is little reason to limit our analysis to propaganda instead of considering all misinformation. Section 2 will present evidence that misinformation influences public opinion, and that even sophisticated consumers have few reasons to believe that they are immune to it. I will discuss instances of intentional misinformation as well as more ambiguous cases that can reasonably be understood as involving the accidental communication of faulty information. Section 3 explains how misinformation poses a challenge to the non-ideal theory of epistemic democracy by discussing the Numbers Trump Ability Theorem discussed by Landemore. Finally, section 4 considers the strengths and weaknesses of democratic deliberation and fact-checking as corrective strategies against misinformation.

Before I begin, two methodological clarifications are in order. First, although the paper will discuss various instances of misinformation, I will concentrate on *online* misinformation (such as “fake news”), which proliferates better and faster than any other kind. My intention in doing so is to contribute to a growing philosophical subfield: the political epistemology of online interaction. Second, my intention is not to offer a general critique of the ideal of epistemic democracy. Epistemic democrats make several idealizing assumptions in order to defend the epistemic value of democratic decision making, but misinformation is a political problem that affects the democracies we currently live in. My interest is in assessing whether the benefits of democratic decision-making discussed by epistemic democrats translate to the real world.<sup>1</sup>

## I. PROPAGANDA AND MISINFORMATION

Like many political buzzwords, such as populism, demagoguery, and ideology, propaganda is notoriously difficult to define. This can be explained by the fact that it mostly functions as a *Kampfbegriff*, that is, a rhetorical weapon rather than a precise analytical tool. If we accuse

---

<sup>1</sup> It is worth noting that my argument does not in any way rely on the claim that ideal political theorizing is not a worthwhile philosophical project.

someone of engaging in propaganda, the implied judgment is usually that this person intends to achieve questionable political ends by misleading the masses.

Stanley attempts a precise definition of propaganda so that we can see it as something other than a vague notion used by political rivals seeking to win a public argument, but it is unclear that he has succeeded in turning the concept into something other than a political cudgel. One can see this in the peculiar specificity of his definitions. His broadest definition of propaganda (2015, xiii) is “the employment of a political ideal against itself” (Stanley 2015, xiii); he further differentiates between “supporting” and “undermining” propaganda. Supporting propaganda is “a contribution to public discourse that is presented as the embodiment of certain ideals, yet is of a kind that tends to increase the realization of those very ideals by either emotional or other nonrational means.” Undermining propaganda is “a contribution to public discourse that is presented as an embodiment of a certain ideal, yet is of a kind that tends to erode those very ideals.” Thus, according to Stanley, a patriotic appeal to a romantic vision of a country’s history is a form of supporting propaganda. When delivering an impassioned nationalistic speech, a politician tries to foster love of country by affecting emotions rather than by presenting arguments. Undermining propaganda is both more nebulous and too narrow. Many typical instances of propaganda do not seem to fit. For instance, Jason Brennan (2017, 8) notes that “instances of wartime propaganda in which the Germans or Japanese are depicted as subhuman animals” do not clearly reference a specific ideal. Brian and Samuel Leiter (2015) remark that “much of the propaganda that mobilized support in America for the 2003 invasion of Iraq . . . involved simple falsehoods or misleading statements” and was not presented as the embodiment of ideals that it implicitly distorted or undermined. As we will see below, the failures of intelligence and communication surrounding the Iraq war are more complicated than appear at first sight, but the authors’ point is that Stanley’s requirement that communicative acts must subvert ideals to qualify as undermining propaganda significantly restricts the extension of the concept, and such a restriction seems arbitrary—perhaps an attempt to equate undermining propaganda with *fascist* propaganda that subverts democratic ideals. A more inclusive definition of propaganda is put forward by Randal Marlin (2013). In his view, propaganda is “the organized attempt, through communication, to affect belief or action or inculcate attitudes in a large audience in ways that circumvent an individual’s adequately informed, rational, reflective judgment.” One advantage of this definition is that it includes communicative acts that we usually count as propaganda even if they are not presented as the embodiment of ideals that they erode. Thus, simple falsehoods and misleading statements intentionally directed at a large audience for political purposes fall within this definition, as they are the result of attempts to bypass individuals’ rational capacities by inciting them to form false beliefs. Thus, Marlin’s definition does not suffer from the arbitrary narrowness of Stanley’s.

Another advantage of Marlin’s definition, however, is that it points us beyond propaganda to the wider phenomenon of misinformation. Marlin’s definition registers the normal assumption that propaganda is the *deliberate* communication of faulty information, but the communication of faulty information is not always the result of a lie; it can be unintentional.<sup>2</sup>

---

<sup>2</sup> One way to accommodate the distinction between deliberate and accidental cases of misinformation would be to distinguish between “disinformation” (intentional cases) and “misinformation” (accidental cases), as is sometimes done in communication studies. I have no objection to this terminology, but I do not find it intuitive and I struggle to see what advantage it has over continuing to use qualifiers.

Misinformation, then, can be defined as the dissemination of false or misleading information, often on a mass scale. Sometimes, as in propaganda, the disseminators communicate what they know (or think) are falsehoods, for any number of political or commercial purposes (not just those singled out by Stanley). But sometimes misinformers may themselves be misinformed, simply because they erroneously believe, and repeat, what is false.<sup>3</sup> When this happens, it can hardly be conceived of as propaganda for it does not involve an *attempt* to bypass individuals' rational capacities, but may simply result from failures in information gathering or communication.

Unintentional misinformation may nevertheless have a negative impact on the quality of political decisions made by the citizenry. My suggestion in the rest of this paper will be that misinformation of both the intentional propagandistic kind *and* the inadvertent kind undermine the epistemic value of democracy. If our concern is the epistemic potential of democratic decision making, therefore, there are few reasons to focus solely on propaganda and neglect accidental misinformation. What is more, many cases of misinformation do not fit squarely on one side of the intentional/unintentional divide, as faulty information is often spread both by individuals who believe that the claims they share with others are false and by individuals who believe such claims to be true. Indeed, misinformation cascades customarily involve individuals who deliberately produce and diffuse false or misleading claims while counting on the fact that it will be relayed by individuals who do *not* perceive it as flawed. This is precisely how fake news—at least that produced by so-called “troll farms”—spreads on social networking sites.

Before examining such cases, however, allow me one more conceptual clarification on the relationship between misinformation and propaganda. I suggested that cases of accidental misinformation do not qualify as propaganda. Yet the reverse is also true: not all instances of propaganda qualify as misinformation. To see this, consider the diverse means used by propagandists. Propaganda sometimes aims to trigger emotions and elicit desires: think, for instance, of World War II “We can do it!” posters, a slogan designed to motivate women to aid in the war effort by replacing male factory workers who were engaged in the military. In other cases, propaganda mostly affects beliefs. By way of example, a poster on which we can read that “Migrants commit more crimes than native citizens”—perhaps set against a background image of migrants waiting at a border crossing—aims to induce the belief that immigration is dangerous. I therefore propose to distinguish between the following three ideal types of propaganda:

\* Affective propaganda: The organized attempt, through communication, to *trigger emotions* in a large audience in ways that circumvent its members' adequately informed, rational, reflective judgment.

\* Conative propaganda: The organized attempt, through communication, to *elicit desires* in a large audience in ways that circumvent its members' adequately informed, rational, reflective judgment.

\* Cognitive propaganda: The organized attempt, through communication, to *shape beliefs* in a large audience in ways that circumvent its members' adequately informed, rational, reflective judgment.

---

<sup>3</sup> Technically speaking, misinformation can happen on a small scale. For instance, I may misinform my wife regarding the hour at which our friend's dinner party will start. In the rest of the discussion, however, my focus is on cases of misinformation directed at a large audience.

I define these three kinds of propaganda as ideal types given that numerous concrete instances of propaganda are not purely affective, conative or cognitive. In fact, they are often all of these at once: a placard on which people from a specific ethnicity are compared to animals or insects, for example, simultaneously aims to trigger disgust, induce the false belief that these people are inferior, and elicit the desire to rid the country of them. Nevertheless, distinguishing between the three ideal types is analytically useful: it helps us scrutinize instances of propaganda from different angles and identify its various goals.

Contrary to those who design fake news, Uncle Sam did not want to disseminate false beliefs: as he made very clear, he wanted you to join the U.S. Army, that is, to induce behavior by eliciting the desire to defend the United States. By way of contrast, misinformation is essentially cognitive: it primarily affects beliefs. Moreover, when it elicits emotions and triggers the desire to act, it usually does so *through* the shaping of one's beliefs. Most cases of what I define as cognitive propaganda involve the communication of faulty or misleading information, and I therefore propose to conceive of cognitive propaganda as the subset of misinformation that involves the intentional deceit (or attempted deceit) of large segments of the population for commercial or political purposes.<sup>4</sup> The majority of instances of misinformation that I will examine fit this category. Yet, as I would now like to demonstrate, cases of misinformation that do *not* unambiguously display an intention to deceive the public permeate democratic life and should not fly under the theoretical radar of political epistemologists.

## II. ASSESSING THE EFFECTIVENESS OF MISINFORMATION

Misinformation would not be a challenge to the epistemic value of democracy if we did not have reasons to believe that it influences the beliefs of those who are exposed to it. Yet recent findings in political science, communication studies, and social psychology suggest that it does.

Consider first the example of the framing of the 2003 invasion of Iraq. Prior to the invasion, members of the Bush administration, including the President and the Vice President, made statements suggesting that weapons of mass destruction (WMD) had been found in Iraq and that Saddam Hussein was working closely with al-Qaeda. Whether they did so because they believed these claims to be true or were attempting to deceive the public for political purposes remains a matter of debate. The Robb-Silverman report, conducted in the wake of the invasion, indicates that the American intelligence community did believe that there were WMD in Iraq, presumably due to failures in intelligence analysis (Jervis 2010, ch. 3). Thus, portraying the Bush administration's statements as an unambiguous case of intentional misinformation amounts to an oversimplification. Regardless of the explanation for the statements, however, they seem to have fostered erroneous beliefs.

In 2003, the Program on International Policy Attitudes conducted a series of polls about misperceptions regarding the war in Iraq. In June 2003, three months after the invasion began, 48

---

<sup>4</sup> Some scholars interested in the study of propaganda may resist my categorization of cognitive propaganda as a subset of misinformation by pointing out that, in some instances of propaganda, propagandists aim to shape the beliefs of individuals without exposing them to faulty or misleading information. Such scholars will be inclined to conceive of the relationship between misinformation and cognitive propaganda as the one that exists between two circles of a Venn diagram rather than—as I propose to do—as two concentric circles. Either way, however, political epistemologists should attend to all cases of misinformation, not solely those that can also be considered cases of cognitive propaganda.

percent of the respondents believed that clear evidence that Saddam Hussein was working closely with al-Qaeda had been found, and 35 percent of the respondents believed that weapons of mass destruction had been found in Iraq.<sup>5</sup> As Steven Kull, Clay Ramsay, and Evan Lewis (2003, 597) note, “among those who did not hold the key false beliefs, only a small minority supported the decision to go to war.” In fact, “in a regression analysis, the presence of misperceptions was the most powerful factor predicting support for the war.”

A second heavily mediatized topic that lends itself to misleading statements is climate change. As in the previous case, faulty information on this subject is not always disseminated by individuals who unmistakably attempt to deceive the public. This can be explained by the fact that the scientific literature on anthropogenic global warming is intricate and is necessarily a matter of estimates with varying ranges of confidence, and that news stories covering it are prone to oversimplification. For instance, news articles from liberal outlets sometimes suggest that most scientists expect catastrophic changes to happen in the next ten years while there is in fact disagreement on this issue. On the other side of the political spectrum, conservative media often disseminate the views that climate change is not occurring or is not anthropogenic despite the scientific consensus that anthropogenic global warming is real. In a recent study on the relationship between media consumption and perceptions of global warming, Jay D. Hmielowski, Lauren Feldman, Teresa A. Myers, Anthony Leiserowitz, and Edward Maibach (2013, 866) found that “conservative media use decreases trust in scientists which, in turn, decreases certainty that global warming is happening.”<sup>6</sup> This is consistent with the finding that “watching Fox News (Feldman et al. 2012; Krosnick and MacInnis 2010), consuming news stories that present evidence questioning the certainty of climate change (Corbett and Durfee 2004), and watching stories that include an interview with a skeptical scientist commenting on global warming (Malka et al. 2009) all decrease belief that global warming is happening and human caused” (Hmielowski et al. 2013).

In these cases, there is room for doubt that the communication of misleading information was or is deliberate. Let us now turn to the effects of a growing kind of misinformation that can more hardly be conceived of as accidental. During the last few years, we have witnessed the rise of fake news, which, borrowing Axel Gelfert’s (2018, 11) definition, we can describe as “the deliberate presentation of false or misleading claims *as news*, where the claims are misleading by

---

<sup>5</sup> It may be possible to explain the widespread belief that Hussein was working with al-Qaeda by the public’s unfamiliarity with Middle Eastern politics, but it is likely that Americans who believed that WMD had been found in Iraq were influenced by the Bush administration’s claims. As Kull, Lewis and Ramsay note, “people do not form misperceptions in a vacuum” (2003, 570).

<sup>6</sup> One methodological concern over studies that associate conservative media use and belief in anthropogenic global warming is that we have evidence that people seek out sources that confirm their opinions (selective exposure). Prima facie, it is therefore possible that lack of belief in global warming could have caused people to use conservative media, not vice versa. Yet, one reason to think that the media have a substantial impact on belief is that modifying the media content to which a group of randomly selected individuals is exposed also modifies their beliefs (on this general phenomenon see Althaus 2003). For example, when Corbett and Durfee (2004) exposed a group of subjects to different versions of a story about a scientific study reporting that the Antarctic Ice Sheet is melting, they found that subjects exposed to the story to which controversy was deliberately added reported a lower level of certainty in global warming than subjects placed in the control condition. This suggests that by repeatedly insisting on the so-called lack of consensus on anthropogenic climate change in the scientific community, conservative broadcasters do foster doubts.

design.” “The originator of an instance of fake news,” he explains, “either intends a specific claim to be misleading in virtue of its specific content, or deliberately deploys a process of news production and presentation that is designed to result in false or misleading claims.” As we have seen, this does not require that all individuals involved in a misinformation cascade have the intention to deceive others. Consider for instance a much-discussed example of fake news: the claim that high officials of the Democratic Party were involved in a sex-trafficking ring. Before becoming an instance of fake news, this conspiracy theory was circulated on a discussion website (reddit.com) by individuals who arguably believed it to be true. It was then transformed into a fake news article and published on a website that specializes in such content (Yournewswire.com; now newspunch.com). The story eventually made its way to social networking sites, and finally went viral. In this instance, a conspiracy theory was exploited and transformed into a fake news article by individuals who hoped to profit from its diffusion; it was then relayed on social networks by users who deemed it sufficiently plausible to be shared. In other cases, fake news is politically motivated and fabricated by individuals who aim to bolster support for a specific cause. Indeed, false claims are often presented in the format of news articles by people who believe them to be false, and who mimic the conventions of traditional media reportage to make them look true (Rini 2017, E-45).<sup>7</sup>

Researchers from the Oxford Internet Institute suggest that fake news is now widespread on social networking sites. Lisa-Maria Neudert (2017, 23) found that “misinformation and junk news content play a substantial role on German social media, accounting for roughly 20 percent of all political news and information on Twitter.” During highly politicized periods such as presidential elections, this 4:1 ratio of political news and information to misinformation can increase to 1:1 (ibid., 17). For their part, Andrew Guess, Brendan Nyhan, and Andrew Reifler (2018) found that approximately one in four Americans visited a fake news website during the 2016 election campaign. A more conservative estimate comes from Allcott and Gentzkow (2017), who found that the average American encountered between one and three stories from known publishers of fake news during the month before the 2016 election.<sup>8</sup>

In the largest study of online misinformation available as of this writing, Soroush Vosoughi, Deb Roy and Sinan Aral (2018, 1146) found that “falsehood diffused significantly farther, faster, deeper and more broadly than the truth in all categories of information,” and that “the effects were more pronounced for false political news” than for other types of fake news. Studying the tweeting and retweeting of 126,000 statements during the entire lifespan of Twitter, they more precisely concluded that the truth rarely diffuses to more than 1000 people while the top 1 percent of fake news “routinely diffused between 1000 and 100,000 people” (ibid., 1148).

---

<sup>7</sup> Is it not possible to conceive of cases in which individuals sincerely believe even the most far-fetched claims to be true, and then choose to present them as news in an attempt to inform the public? It is not impossible, and it is important to acknowledge that intentions are notoriously difficult to track. But when a fake news article can be traced back to an organization that repeatedly diffuses unsubstantiated claims and that has a commercial or political interest in doing so—think for instance of the Russia-based Internet Research Agency—a reasonable conclusion to draw is that such an organization is knowingly deceiving the public. I thank Jeffrey Friedman for pressing me on this point, and more generally for inciting me to draw the distinction between intentional and unintentional cases of misinformation.

<sup>8</sup> This estimate is conservative because the study only tracked 156 fake news stories (Allcott and Gentzkow 2017). It also did not measure exposure to misinformation on mobile devices or cases in which individuals were exposed to misinformation on their newsfeed without visiting the fake news website (i.e. without clicking on the link).

Moreover, it takes the truth about six times as long as falsehood to reach 1500 people, and falsehoods are 70 percent more likely to be retweeted than the truth. Like Neudert, these researchers found that the total number of false rumors peaks during highly politicized periods such as U.S. presidential elections and the Russian annexation of Crimea.

How is it possible that masses of people are often led to believe that which is simply not the case? Fake news articles often contain false statements so far-fetched that we might plausibly suppose that individuals must have the critical abilities necessary to identify them and avoid letting falsehoods influence their opinions. Should we not believe, then, that individuals are immune to misinformation? Such an assumption needs to be rethought. Indeed, at least five psychological findings suggest that misinformation can influence people's beliefs. First, "when people know little about a source, they treat information from that source as credible" (Rapp 2016, 284). In fact, it is very useful for us to do so, as we often gain valuable insights from strangers by giving them our epistemic trust; constantly doubting the veracity of others' claims would hinder the achievement of our goals. Those who deliberately engage in online misinformation exploit this fact by creating unknown websites the content and design of which imitate those of traditional media outlets. Think, for instance, of the now-defunct *Denver Guardian*, a fake news website that was virtually indistinguishable from the real newspaper *The Denver Post* (for non-Coloradans, at least).<sup>9</sup>

Second, psychologists have long been interested in what they call the "truth effect," which results from repeated exposure to a piece of information or misinformation. When presented with the same piece of mis/information multiple times, we are more likely to believe it.<sup>10</sup> As Neil Levy notes, repetition increases processing fluency—the subjective ease of information processing—and people tend to accept claims that are processed fluently. In fact, Levy (2017, 29) argues that fluency effects may lead us to eventually accept claims we thought of as false:

Repetition of a claim powerfully affects fluency of processing. This effect may lead to the agent accepting the original claim, when she has forgotten its source. Even when repetition is in the service of debunking a claim, it may result in higher levels of acceptance by promoting processing fluency.

This is of interest given that we are repeatedly exposed to misinformation, either because our friends and families post and share it online or because we read genuine newspaper articles that try to debunk them.

Third, people "rely on fiction as a source of information, even when fiction contradicts relatively well-known facts about the world" (Marsh and Fazio 2006). For instance, when Deborah Prentice, Richard Gerrig, and Daniel Baillis (1997) asked participants in an experiment to read two versions of a fictional story—one of which contained the claim that mental illness is contagious while the other did not—they found that individuals who read the stories associating mental illness with contagion were significantly more likely to believe that mental illness really is contagious. Moreover, acceptance of false claims does not decline when participants are warned that the stories presented may contain false information (Marsh and Fazio 2006). As it

---

<sup>9</sup> The man behind the *Denver Guardian* publicly admitted that he created and diffused fake news article both for profit and for his own amusement (Sydell 2016).

<sup>10</sup> For a review, see Dechêne, Stahl, Hansen, and Wänke 2010.



turns out, the intuitive idea according to which a fake news article will not influence the beliefs of a reader who knows that he is reading fake news is not supported by empirical evidence.

Fourth, people sometimes fail to recall the source of the information they acquire. People who are first exposed to a claim while reading a fake news article may eventually misattribute this claim to a reliable source or to common knowledge (Marsh, Cantor and Brashier 2016; Rapp 2016). Such psychological findings lead Levy to reject what he defines as the naïve view of belief and its role in behavior, that is, the view that

mental representations are *reliably* and *enduringly* categorized into kinds: beliefs, desires, fantasies and fictions, and that we automatically or easily reclassify them given sufficient reason to do so. On this picture, fake news is a problem when it results in representations that are categorized as beliefs. That problem is averted by ensuring that the representations we form as we consumer fake news are not wrongly categorized. (Levy 2017, 22).

Contrary to this view, however, the available evidence suggests that we often unknowingly and unintentionally reclassify fantasies and fictions as justified true beliefs.

Finally, a fifth finding of interest is that fake news articles arouse surprise and disgust more strongly than does true information (Vosoughi et al. 2018, 1150), which makes them more likely to be shared and become part of our worldview. Thus, as Kim Peters, Yoshihisa Lashima, and Anna Clark (2009, 2007) note, “emotional social events (particularly those that arouse disgust and happiness) are likely to become part of a society’s social beliefs.”<sup>11</sup>

Contemporary psychological findings thus suggest that repeatedly reading or hearing that Hillary Clinton sold arms to a terrorist organization or that the pope endorsed Donald Trump during the 2016 U.S. presidential campaign will affect cognition, even in cognitively sophisticated individuals who know that the sources to which they are exposed are communicating false claims.

In summary, there is empirical evidence that misinformation effective, and that confidence in our ability to immunize ourselves from its effects is unwarranted. Still, some epistemic democrats may accept these claims while rejecting the idea that misinformation endangers collective wisdom.

### III. EPISTEMIC EFFECTS OF MISINFORMATION IN REAL DEMOCRACIES

Epistemic democrats argue that the value of democracy does not entirely depend on the intrinsic fairness of democratic procedures. If democracy is to be considered legitimate, it must also be the case that such procedures yield good normative outcomes (Estlund and Landemore 2018). A weak version of the epistemic argument for democracy is defended by David Estlund, who argues that democratic decision-making procedures are better than random selection procedures at avoiding primary bads.<sup>12</sup> A stronger version of the epistemic argument is defended by

---

<sup>11</sup> Note, however, that this fifth finding supports the narrower claim that *sensationalistic* misinformation (rather than misinformation as such) is effective.

<sup>12</sup> Primary bads include “war, famine, economic collapse, political collapse, epidemic, and genocide” (Estlund 2008, 163). Epistemic democrats thus make the moral assumption that some outcomes can be deemed good or bad

Landemore (2013, 3), who argues that a “democratic decision procedure is likely to be a better decision procedure than any non-democratic decision procedures, such as a council of experts or a benevolent dictator.” Although Landemore is primarily interested in the “ideal of democratic decision-making” (ibid., 9), she also believes that there is value in discussing whether the theoretical assumptions behind the various theorems she uses to defend the epistemic value of democratic decision-making plausibly translate to the real world.

In this section, I argue that there are significant reasons to doubt that the most important theorem on which Landemore relies to establish the epistemic value of democracy—the Numbers Trump Ability Theorem—applies to real democracies in the modern world. More specifically, my contention is that misinformation creates a gap between the degree of political competence people must have if the mechanisms of collective wisdom are to function and the epistemic situation of our actual compatriots. The conclusion of this argument is not that epistemic democracy is valueless as an ideal, but that misinformation currently prevents the citizens of mass democracies from reaping the benefits of the “wisdom of crowds.”<sup>13</sup> Landemore’s Numbers Trump Ability Theorem draws on Scott Page’s Diversity Trumps Ability Theorem. Relying on computer models consisting of simulated agents, Page (2007, 162) demonstrates that “a randomly selected collection of problem solvers outperforms a collection of the best individual problem solvers.” This is explained by the fact that the group of randomly selected problem solvers has an epistemic advantage over the collection of the best individual problem solvers: cognitive diversity. To solve problems in an efficient manner, it is better to give decisional powers to a group of cognitively diverse people than to a group of very smart people who think alike. According to Landemore (2013, 104), we can draw democratic insights from Page’s demonstration. Indeed, democratic modes of decision-making such as collective deliberation are eminently egalitarian and inclusive, so there is reason to believe that they effectively foster cognitive diversity. “To the extent that cognitive diversity is a key ingredient of collective intelligence,” Landemore writes, “and specifically one that matters more than average individual ability, the more inclusive the deliberation process is, the smarter the solutions

---

according to a normative criterion that is independent from the collective will of the people. I share this assumption. For instance, we can say that genocide is bad even if a large majority of voters support it.

<sup>13</sup> Although my argument is an exercise in non-ideal theorizing about democracy, it is worth noting that the central idealizing assumptions made by epistemic democrats may not immunize their theories against the problem posed by misinformation. For instance, Landemore (2013, 11 and 197) assumes that citizens and rulers are motivated by a concern for the public good and actively seek to promote what they conceive of as the interests of “the entire community, as opposed to the interests of the rulers themselves or a subgroup within society.” A second idealizing assumption central to her theory is that democratic deliberation has epistemic merits only when individuals “properly deliberate,” that is, when all deliberating parties engage in “genuine consideration of arguments for and against something” (ibid., 138). This may seem to preclude people from misinforming each other—but only if those who engage in misinformation are motivated by a concern for the interests of a subgroup within society. When pundits deliberately lie on social media during electoral campaigns, they may be motivated by the consequentialist reasoning that this will ultimately serve the common good: it may help the political party they consider to be truly promoting their country’s interest. As for the assumption that deliberation yields epistemic benefits only when people properly deliberate, it certainly precludes us from conceiving of deliberative processes *during which* individuals misinform each other. Yet misinformation not only affects cognition during deliberation, but also *before* individuals begin to deliberate. Indeed, a deliberative assembly may be composed of well-intentioned individuals who properly deliberate, but who have been misinformed in the past. When this is the case, they risk making political decisions based on false beliefs and misinforming others unwittingly.

resulting from it should be, overall” (ibid., 104). In other words, numbers can realistically function as a proxy for diversity, and the Diversity Trumps Ability Theorem can therefore be generalized into a Numbers Trump Ability Theorem. Democratic masses are more likely to make normatively desirable decisions than are small groups of good problem solvers, and this gives us reason to think that democracy has greater epistemic value than other modes of political decision making, such as the rule of experts.

However, certain conditions must be met for the Diversity Trumps Ability Theorem to hold, as both Page and Landemore admit. More precisely, it must be the case that (1) “the problem must be difficult enough, since we do not need a group to solve easy problems”; (2) “all problem solvers are relatively smart (or not too dumb)”; (3) “the participants think very differently, even though the best solution must be obvious to all of them when they are made to think of it”; and (4) “the initial population from which the problem solvers are picked must be large and the collection of problem solvers working together must contain more than a handful of problem solvers” (ibid.). Critics have argued that Landemore is overly optimistic when affirming that these conditions are not unduly demanding (Ancell 2017). Even Page seems more pessimistic than Landemore about the epistemic merits of diversity. In the prologue of *The Difference* (2007, xxix), he writes:

Of course, difference does not magically translate into benefits. My claims that diversity produces benefits rest on conditions. These conditions require, among other things, that diversity is relevant—we cannot expect that adding a poet to a medical research team would enable them to find a cure for the common cold. . . . Diversity, like everything else (excepting, of course, moderation), has its limits.

A certain level of individual epistemic competence is therefore required so that inclusive democratic deliberation can bear fruit.

Yet misinformation frequently impedes democratic citizens from reaching that level. A group of citizens who believe that vaccines cause autism, for instance, is not likely to propose judicious public health policies to fight measles outbreaks. Many instances of misinformation disseminate false beliefs that impair individuals’ ability to make sound normative judgments: if one faces the task of determining whether the government should try to reduce carbon emissions, holding the true belief that most climate scientists conceive of global warming as serious and anthropogenic will be of great help. If one must decide whether to invade a foreign country, holding the false belief that this country’s government hides weapons of mass destruction will undermine her reasoning. One influential study that supports the conclusion that holding false beliefs influences political reasoning focuses on Americans’ beliefs about the percentage of families who are on welfare, the average annual benefit amount for a recipient family, the proportion of the federal budget that welfare absorbs, and the percentage of recipient families who are African-American (Kuklinski et al. 2000). On none of these items did a majority of participants, or close to it, get the facts right. In fact, the proportion getting an individual fact wrong ranged “from two-thirds on the percentage of all welfare who are African-American to a striking 90 percent on the percentage of the federal budget that goes to welfare.” What is more, those who held the least-accurate beliefs expressed the highest confidence in them. For instance, “47% of those who estimated the proportion of American families on welfare correctly (at 7%)

said they were very or fairly confident, while 74% of those who grossly overestimated the figure (at 25%) did.” Relying on simulations, James Kuklinski, Paul Quirk, Jennifer Jerit, David Schwieder, and Robert Rich also found that misinformation skews political preferences in both pro- and anti-welfare ways, at least if we compare groups of subjects who are highly misinformed with subjects that are only moderately so. As the authors remark, “it is one thing to find misinformed citizens, quite another to show that misinformation has an effect on the citizenry’s collective voice” (ibid., 805). Their findings suggest the latter.<sup>14</sup>

The challenge posed by organized attempts to disseminate false beliefs is also distinct from one of the most frequent objections to the idea that democratic decision-making can yield good outcomes through majority rule: the argument from ignorance. Bryan Caplan (2007), Ilya Somin (2016), and Jason Brennan (2016) argue that a wide array of empirical studies demonstrate that citizens lack the political knowledge necessary to make informed political decisions. Most Americans do not know which party controls Congress, many cannot identify any congressional candidates in their district, and 40 percent of them do not know whom the United States fought in the Second World War II (Brennan 2016, 25-26). Yet, as the Oracle of Delphi taught Socrates, being ignorant is not equivalent to holding a false belief; one may be wise just by knowing that one does not know. A voter who is consciously ignorant of the two presidential candidates’ views on foreign policy may deliberately refrain from making her choice on that basis and rely instead on knowledge she does possess. By way of contrast, a voter who cares greatly about foreign policy but has been misinformed—that is, one who holds false beliefs—about the candidates’ positions on that matter is more likely to make a poorly justified choice. Falsely believing that x is not that case is an epistemic state distinct from not knowing whether x is the case, and it is more dangerous from a normative point of view. To illustrate, consider the example of Paul and Mary, two American citizens who both care greatly about immigration. Mary knows little about the actual effects of immigration on the economy and crime rates, so she decides to support the candidate who is most committed to fight the effects of climate change on the environment, an issue about which she also greatly cares. Paul also cares greatly about immigration, but has been misinformed about its effects. He has read fake news that led him to believe that a majority of immigrants to the United States commit crimes, and he therefore decides to vote for the candidate who proposes to close the borders.<sup>15</sup> Mary is in a state of

---

<sup>14</sup> Other studies suggesting that holding false beliefs influences political and voting preferences include Hochschild and Einstein 2014 and Reedy et. al. 2014. As is often the case when considering empirical evidence, it remains difficult to draw definitive conclusions from such studies. For this reason, my suggestion in this section and the previous one is not that we should draw such conclusions, but, more moderately, that findings in social and political psychology give us reasons to believe that misinformation will be an obstacle to the translation of the ideal theory of democracy to the real world. At the very least, it is a problem that should be discussed by epistemic democrats interested in non-ideal theory.

<sup>15</sup> Of course, political judgments may be overdetermined. For instance, someone might support greater immigration because (a) she believes that migrants have a positive impact on her country and (b) she believes that migrants have a moral right to emigrate to her country. In this case, coming to falsely believe that migrants committed crimes that were in fact fabricated will not modify her overall judgment on immigration if she considers the reasons that stem from (b) as overriding. But false beliefs may also give rise to (internal) reasons that are sufficiently strong as to modify an individual’s overall judgment. As we have seen, Kull, Ramsay, and Lewis’s study suggests that the false belief that WMD had been found in Iraq boosted support for the war. A study by Richard Gunther, Paul Beck, and Erik Nisbet (2018), still in its initial stages, appears to suggest that Obama voters who held false negative beliefs about Hillary Clinton were 3.9 times more likely to vote for Trump than those who did not. This is hardly surprising:

Socratic ignorance. She knows that she does not know a lot about immigration, so she bases her vote on other considerations. Paul is not in an epistemic state which allows him to do so, as he assumes that the unsubstantiated beliefs he holds regarding immigration and crime amount to knowledge. As Kuklinski and his associates point out (2000, 799), “although factual inaccuracy is troublesome, it is the ‘I know I’m right’ syndrome that poses the potentially formidable problems.” Individuals are not merely in the dark, they are often wrong-headed and yet confident that they are knowledgeable.

To temper the claim that *ignorance* impairs political decision-making, a common strategy is to dissociate political competence from factual knowledge and argue that knowledge of objective facts is an elitist measure of political knowledge that is irrelevant to the ability to cast a politically competent vote (Lupia 2006). There is undoubtedly some truth to this, as many surveys used to measure political knowledge focus on such questions as how many senators each state has or what the names of Supreme Court Justices are—knowledge that only elites or political junkies may know. Nonetheless, some factual knowledge is required for one to form sound political judgments. Falsely believing that the most migrants commit crimes will hamper our reasoning if the task at hand is to determine if we should vote for the candidate who wants to close borders or for the candidate who argues that we ought to accept more refugees.<sup>16</sup> It may be true that Caplan, Somin, and Brennan overestimate the amount of knowledge necessary to make sound political judgments, but conversely, epistemic democrats may underestimate the impact of *either* ignorance or false beliefs. Ignorance of important information can lead to ill-considered political decisions, and so, a fortiori, can important nuggets of misinformation.<sup>17</sup>

#### IV. INEFFECTUAL DELIBERATION

Although I have attempted to paint a realistic portrait of the effects of misinformation, it might be thought that I have been unrealistic in treating the victims of misinformation as isolated individuals—isolated, that is, from the corrective effects of deliberation.

James S. Fishkin and Robert C. Luskin have found that participants in deliberative polling experiments gain significant knowledge after spending a weekend discussing a preselected political issue with other citizens and consulting with experts on that matter. In a

---

falsely believing that a presidential candidate committed immoral actions provides one with strong reasons to vote for her opponent. I thank Emily Sullivan for drawing my attention to this point.

<sup>16</sup> The pernicious effects of misinformation extend beyond the decision to vote for candidate x. Given the importance of surveys in the political process of Western democracies, the prevalence of widely held false beliefs may influence the very political options that are available to citizens. If a large segment of the population has a negative view of migrants because of fake news that portrays them in an unflattering light, all political parties are likely to take a tougher stance on immigration to cater to a wide range of voters, at least in countries with majoritarian electoral systems. I owe this point to Zoe Phillips Williams.

<sup>17</sup> In this paper, I limit myself to an assessment of the political consequences of misinformation. More precisely, my suggestion is that false beliefs lead to poorly justified political decisions, and that poorly justified political decisions risk producing bad outcomes. I share this assumption with political scientists concerned about the problem of political ignorance, who also suppose that it is better for a democracy to have citizens who make adequately justified decisions than one that does not. That said, one could extend the analysis of misinformation by arguing that it is undesirable for individuals to hold false beliefs, regardless of the societal consequences. From an instrumental point of view, people need to construct a reliable picture of the world in order to effectively further their interests, such that misinformation about personal matters may be disastrous. From a perfectionist point of view, one could also argue that knowledge and truth are essential components of a good human life.

1999 deliberative poll on whether Australia should become a republic, “the percentage knowing that ‘under the referendum proposal’ the prime minister ‘could remove the president at any time but must later obtain approval from the house of representative’ rose from 16 to 73%” (Fishkin and Luskin 2005, 291). Insofar as deliberation in the real world mimics the effects of deliberative mini-publics, or insofar as the results of mini-publics could be broadcast to nonparticipants, thereby entering into their deliberations, one might entertain the hope that these effects could lead to a better-informed public.

However, studies of deliberative mini-publics focus on knowledge gains, not on the correction of false beliefs, and there is evidence that convincing individuals to correct their erroneous beliefs is significantly more difficult than teaching them facts that do not contradict the beliefs they already hold. As Levy (2017, 27) notes, “the phenomenon of belief perseverance has long been known to psychologists,” and “corrections rarely if ever eliminate reliance on misinformation.” For instance, Lee Ross, Mark Lepper, and Michael Hubbard (1975) found that when people are asked to perform a task and are given positive feedback on their performance, they do not stop believing that they are better than average at the task when told afterwards that the given feedback was unrelated to their performance. Jonas De Keersmaecker and Arne Roets (2017, 107) applied these insights to the study of fake news and consistently found that the false beliefs they cause prove especially hard to correct for people with low cognitive abilities, even when the false information is “explicitly and unambiguously disconfirmed.”

Given that Landemore’s argument is comparative, it is worth asking if these experimental results strengthen the epistocrat’s case in favor of non-democratic decision procedures. My answer to this question is cautious. Epistocrats can rejoice in De Keersmaecker and Roets’s (2017) finding that false beliefs prove especially hard to correct for people with low cognitive abilities, as it suggests that experts—who are typically individuals with high cognitive abilities—are more likely than lay citizens to be receptive to corrections when their beliefs are contradicted by evidence. Let me nonetheless add two caveats. The first is that there is no reason to believe that experts are immune to misinformation: even if they accept corrections more easily, they may form false beliefs to the same extent as people with low-cognitive abilities, for instance by reclassifying fantasies and fictions as beliefs. The second is that De Keersmaecker and Roets’s study focuses on individuals with high and low cognitive abilities as opposed to groups composed of such individuals, and more work is needed to verify or disprove the claim that a council of experts outperforms diverse democratic assemblies in the correction of false beliefs.

In fact, while it is true that misperceptions are difficult to correct, we should not conclude that all kinds of democratic deliberation are powerless against misinformation. Instead, studies more moderately suggest that certain types of deliberation will more effectively favor the correction of false beliefs than others. When citizens are exposed to information from an individual they do not perceive as ideologically opposed to their worldview, they may engage in epistemic deference, that is, accept this person’s speech as authoritative. Unsurprisingly, individuals tend to trust and respect sources with whom they share an ideological affiliation. Thus, Ifat Maoz, Andrew Ward, Michael Katz, and Lee Ross (2002, 515) found that “Israeli Jews evaluated an actual Israeli-authored peace plan less favorably when it was attributed to the Palestinians than when it was attributed to their own government.” Similarly, Geoffrey L. Cohen (2003) found that Democrats and Republicans assigned greater value to proposals for welfare policies when they were presented as initiatives from their own party. More surprisingly, a recent

study found that “statements from an unlikely source—a person who makes proclamations that run contrary to their personal and political interests”—can increase citizens’ willingness to correct misperceptions “regardless of their own political predilections” (Berinsky 2017, 241). When presented with false claims about the Affordable Care Act (i.e., that it included procedures to withhold care from certain citizens) and with subsequent corrections of these claims, Republicans accepted corrections from other Republicans at a higher rate than any other types of corrections, but corrections from nonpartisan sources (in this case, from the American Medical Association and the American Association of Retired Persons) were also effective. This suggests that collective deliberation in the presence of experts who are perceived to be unbiased can lead to the successful correction of false beliefs.

Thus, deliberative democracy may be a powerful tool against misinformation, at least if the impartiality condition is met. The inclusion of fact-checking in deliberation may also help. A study by **Adrien** Friggeri, Lada Adamic, Dean Eckles and Justin Cheng (2014, 106) showed that reshares of false rumors were 4.4 times as likely to be deleted when debunked than when not. However, fact-checks have their own problems. First, the consumption of fact-checks is concentrated among non-fake news consumers, so they often do not reach their intended audience (Guess, Nyhan and Reifler, 2018). Second, and perhaps most importantly, not all misleading statements can simply be fact-checked.

In many instances, fake news articles contain simple factual assertions that can be shown to be false—or at least unsubstantiated—by journalists. After investigation, there was no evidence that sexual crimes were being committed in the basement of Comet Ping Pong in Washington, D.C., by members of the Democratic Party; in fact, the restaurant had no basement. Yet, many cases of misinformation contain complex statements that are much harder to verify, either because this would require prolonged scientific inquiry or because the truth value of such statements is a matter of interpretation. As Walter Lippmann argued in *Public Opinion* ([1922] 1997, 158), we are frequently exposed to “forms of persuasion that we cannot verify.” Consider predictive statements. NASA currently estimates that the sea level will rise 26 inches by 2100 as a result of climate change. Could we fact-check a fabricated news story that claims that it will not, or that it will only rise by 13 inches? What could be fact-checked is whether such a claim can be attributed to NASA, but not the prediction itself, as only time will tell whether NASA’s prediction or the fabricated one is the most accurate. Second, claims that contain ambiguous terms are often difficult to assess, as judgments regarding their accuracy depend on subjective interpretations. Is it true, for instance, that Donald Trump “encouraged” his supporters to commit violent actions against protestors in the context of political rallies? This depends on what counts as encouraging, which is a matter for debate. Third, misleading news stories often combine unsubstantiated statements with claims that the available evidence supports as true (Uscinski and Butler 2013).<sup>18</sup> In such cases, categorizing an entire news article as “true” or “false,” as fact checkers do, is itself misleading.

It seems to me, then, that at this early stage of research, we cannot rely on the hope that deliberation and fact-checking alone are likely to correct the problem of misinformation.

---

<sup>18</sup> In fact, combining falsehoods with a kernel of truth to make the former appear more convincing is a propaganda technique that was widely used by misinformation officers during the Cold War.

## REFERENCES

- Allcott, Hunt, and Matthew Gentzkow. 2017. "Social Media and Fake News in the 2016 Election." *Journal of Economic Perspectives* 31(2): 211-236.
- Althaus, Scott L. 2003. *Collective Preferences in Democratic Politics*. New York: Cambridge University Press.
- Ansell, Aaron. 2017. "Democracy Isn't *That* Smart (But We Can Make It Smarter): On Landemore's *Democratic Reason*." *Episteme* 14(2): 161-75.
- Aristotle. 1984. *Complete Works*. Princeton: Princeton University Press.
- Berinsky, Adam J. 2017. "Rumors and Health Care Reforms: Experiments in Political Misinformation." *British Journal of Political Science* 47(2): 241-262.
- Brennan, Jason. 2016. *Against Democracy*. Princeton: Princeton University Press.
- Brennan, Jason. 2017. "Propaganda about Propaganda." *Critical Review* 29(1): 1-15.
- Caplan, Bryan. 2008. *The Myth of the Rational Voter: Why Democracies Choose Bad Policies*. Princeton: Princeton University Press.
- Cohen, Geoffrey L. 2003. "Party Over Policy: The Dominating Impact of Group Influence on Political Beliefs." *Journal of Personality and Social Psychology* 85(5): 808-22.
- Corbett, Julia B., and Jessica L. Durfee. 2004. "Testing Public (Un)certainly of Science: Media Representations of Global Warming." *Science Communication* 26(2): 129-51.
- Dechêne, Alice, Christoph Stahl, Joachim Hansen, and Michaela Wänke. 2010. "The Truth about the Truth: A Meta-Analytic Review of the Truth Effect." *Personality and Social Psychology Review* 14(2): 238-57.
- De Keersmaecker, Jonas, and Arne Roets. 2017. "'Fake News': Incorrect, but Hard to Correct. The Role of Cognitive Ability on the Impact of False Information on Social Impressions." *Intelligence* 65: 107-110.
- Estlund, David M. 2008. *Democratic Authority: A Philosophical Framework*. Princeton: Princeton University Press.
- Estlund, David, and Hélène Landemore. 2018. "The Epistemic Value of Democratic Deliberation." In *The Oxford Handbook of Deliberative Democracy*, ed. André Bächtiger, John S. Dryzek, Jane Mansbridge, and Mark E. Warren. New York: Oxford University Press.
- Feldman, Lauren, Edward W. Maibach, Connie Roser-Renouf, and Anthony Leiserowitz. 2012. "Climate on Cable: The Nature and Impact of Global Warming Coverage on Fox News, CNN, and MSNBC." *International Journal of Press/Politics* 17(1): 3-31.
- Fishkin, James S., and Robert C. Luskin. 2005. "Experimenting with a Democratic Ideal: Deliberative Polling and Public Opinion." *Acta Politica* 40(3): 283-98.
- Friggeri, Adrien, Lada A. Adamic, Dean Eckles, and Justin Cheng. 2014. "Rumor Cascades." *Proceedings of the International Conference on Weblogs and Social Media*. Association for the Advancement of Artificial Intelligence (AAAI).
- Hmielowski, Jay D., Lauren Feldman, Teresa A. Myers, Anthony Leiserowitz, and Edward Maibach. 2014. "An Attack on Science? Media Use, Trust in Scientists, and Perceptions of Global Warming." *Public Understanding of Science* 23(7): 866-883.
- Hochschild, Jennifer, and Katherine L. Einstein. 2015. "'It Isn't What We Don't Know that Gives Us Trouble, It's What We Know That Ain't So: Misinformation and Democratic Politics." *British Journal of Political Science* 45(3): 467-75.
- Gelfert, Axel. 2018. "Fake News: A Definition." *Informal Logic* 38(1): 84-117.
- Guess, Andrew, Brendan Nyhan, and Jason Reifler. 2018. "Selective Exposure to Misinformation: Evidence from the Consumption of Fake News during the 2016 U.S. Presidential Campaign." Manuscript, European Research Council.



- Gunther, Richard, Paul A. Beck, and Erik C. Nisbet. 2018. "Fake News Did Have a Significant Impact on the Vote in the 2016 Election." Working Paper, Department of Political Science, Ohio State University.
- Jervis, Robert. 2010. *Why Intelligence Fails: Lessons from the Iranian Revolution and the Iraq War*. Ithaca: Cornell University Press.
- Krosnick, Jon A., and Bo MacInnis. 2010. "Frequent Viewers of Fox News Are Less Likely to Accept to Accept Scientists' Views of Global Warming." Technical Paper. Stanford University, Stanford, CA.
- Kuklinski, James H., Paul J. Quirk, Jennifer Jerit, David Schwieder, and Robert F. Rich. 2000. "Misinformation and the Currency of Democratic Citizenship." *The Journal of Politics* 62(3): 790-816.
- Kull, Steven, Clay Ramsay, and Evan Lewis. 2003. "Misperceptions, the Media and the Iraq War." *Political Science Quarterly* 118(4): 569-98.
- Landemore, Hélène. 2012. "Democratic Reason: The Mechanisms of Collective Intelligence in Politics." In *Collective Wisdom. Principles and Mechanisms*, ed. Jon Elster and Hélène Landemore. C Cambridge: Cambridge University Press.
- Landemore, Hélène. 2013. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton: Princeton University Press.
- Leiter, Brian, and Samuel Leiter. 2015. "Not Your Grandfather's Propaganda." *The New Rambler: An Online Review of Books*, October 12.
- Levy, Neil. 2017. "The Bad News about Fake News." *Social Epistemology Review & Reply Collective* 6(8): 20-36.
- Lippmann, Walter. [1922] 1997. *Public Opinion*. New York: Free Press.
- Lupia, Arthur. 2006. "How Elitism Undermines the Study of Voter Competence." *Critical Review* 18(1-3): 217-32.
- Malka, Ariel, Jon A. Krosnick, Matthew Debell, Josh Pasek, and Daniel Schneider. 2009. "Featuring Skeptics in News Media Stories About Global Warming Reduces Public Beliefs in the Seriousness of Global Warming." Technical Paper. Stanford University, Stanford, CA.
- Maoz, Ifat, Andrew Ward, Michael Katz, and Lee Ross. 2002. "Reactive Devaluation of an 'Israeli' vs. 'Palestinian' Peace Proposal." *Journal of Conflict Resolution* 46(4): 515-46.
- Marlin, Randal. [2002] 2013. *Propaganda and the Ethics of Persuasion*. Peterborough: Broadview Press.
- Marsh, Elizabeth J., Allison D. Cantor, and Nadia M. Brashier. 2016. "Believing that Humans Swallow Spiders in Their Sleep: False Beliefs as Side-Effects of the Processes that Support Accurate Knowledge." *Psychology of Learning and Motivation* 64: 93-132.
- Marsh, Elizabeth J., and Lisa K. Fazio. 2006. "Learning Errors from Fiction: Difficulty in Reducing Reliance on Fictional Stories." *Memory & Cognition* 34(5): 1140-49.
- Neudert, Lisa-Maria. 2017. "Computational Propaganda in Germany: A Cautionary Tale." Working Paper No. 2017.7. Computational Propaganda Research Project, Oxford Internet Institute.
- Page, Scott E. 2007. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton: Princeton University Press.
- Plato. 1977. *Complete Works*. Indianapolis: Hackett.
- Peters, Kim, Yoshihisa Lashima, and Anna Clark. 2009. "Talking About Others: Emotionality and the Dissemination of Social Information." *European Journal of Social Psychology* 39(2): 207-22.
- Prentice, Deborah A., Richard J. Gerrig, and Daniel S. Bailis. 1997. "What Readers Bring to the Processing of Fictional Texts." *Psychonomic Bulletin & Review* 4(3): 416-20.
- Rapp, David N. 2016. "The Consequences of Reading Inaccurate Information." *Current Directions in Psychological Science* 25(4): 281-85.
- Reedy, Justin, Chris Wells, and John Gastil. 2014. "How Voters Become Misinformed: An Investigation

- of the Emergence and Consequences of False Factual Beliefs.” *Social Science Quarterly* 95(5): 1399-1418.
- Rini, Regina. 2017. “Fake News and Partisan Epistemology.” *Kennedy Institute of Ethics Journal* 27: 43-64.
- Ross, Lee, Mark R. Lepper, and Michael Hubbard. 1975. “Perseverance in Self-Perceptions and Social Perception: Biased Attributional Processing in the Debriefing Paradigm.” *Journal of Personality and Social Psychology* 32(5): 880-92.
- Somin, Ilya. 2016. *Democracy and Political Ignorance*. Stanford: Stanford University Press.
- Stanley, Jason. 2015. *How Propaganda Works*. Princeton: Princeton University Press.
- Sydell, Laura. 2016. “We Tracked Down a Fake-News Creator in the Suburbs. Here’s What We Learned.” *New Public Radio*, 23 November 2016.
- Uscinski, Joseph E., and Ryden W. Butler. 2013. “The Epistemology of Fact-Checking.” *Critical Review* 25(2): 162-80.
- Vosoughi, Soroush, Deb Roy, and Sinan Aral. 2018. “The Spread of True and Fake News Online.” *Science* 359(6380): 1146-51.