



Commentary

The “Echo Chamber” Distraction: Disinformation Campaigns are the Problem, Not Audience Fragmentation



R. Kelly Garrett*

Ohio State University, United States

The importance of the arguments made in “Beyond misinformation” (Lewandowsky, Ecker, & Cook, 2017) is difficult to underestimate. Recognizing that the current crisis of faith in empirical evidence and in the value of expertise has roots that reach far beyond individual-level psychological processes is a crucial step in countering it. As the authors note, there are a host of social, technological, and economic factors that contribute to the situation we face today, and accounting for these interdependent forces will enable stakeholders, including scientists, journalists, political elites, and citizens, to respond more effectively.

Further, the authors clearly articulate why scholars must engage in politically charged debates. The argument that political motivations are driving the emergence of a “post-truth” world has ample precedent: rumors and lies have been used to shape public opinion throughout human history (Allport & Postman, 1965 [1947]; Jamieson, 2015; Knapp, 1944; Mara, 2008; Shibutani, 1966). What is perhaps unique to the present situation is the willingness of political actors to promote doubt as to whether truth is ultimately knowable, whether empirical evidence is important, and whether the fourth estate has value. Undermining public confidence in the institutions that produce and disseminate knowledge is a threat to which scientists must respond.

The primary goal of this response, however, is not to underscore the article’s insights. Those contributions speak for themselves. Lewandowsky, Ecker, and Cook (2017) cover considerable intellectual territory, which necessarily requires brevity in their treatment of complex issues. I aim to explore critically an area that merits additional attention: namely, the authors’ characterization of the online information environment.

The risk is that some readers might misinterpret the authors’ claims about the role of technology in a substantively important way. There is strong empirical evidence that most individuals encounter a range of political viewpoints when consuming news, a fact which is potentially obscured by references to “echo chambers” and “filter bubbles.” This has important implications for the strategies used to counter misinformation. In the absence of echo chambers, promoting contact with belief-challenging corrections is insufficient. Effectively responding to disinformation campaigns requires that we find ways to undermine beliefs that persevere *despite* encounters with counter-evidence.

Diversity of Online Political Information Exposure

In their discussion of the changing media landscape, the authors briefly allude to echo chambers, where “most available information conforms to pre-existing attitudes and biases” (Lewandowsky et al., 2017, p. 353). Although the authors are careful not to assert that this is what most people experience when they go online, readers may nevertheless assume that it is. Yet there is ample evidence that echo chambers are not a typical part of Internet users’ experience. Numerous analyses of large-scale observational data indicate that online news consumers do not systematically avoid exposure to content with which they would be expected to disagree.

Two important studies help illustrate this point. First, a three-month study of 50,000 online news users in the U.S. found that news audience fragmentation was quite low. Specifically, the authors found that, on average, the distance between the ideologies represented in the news diets of a random pair of consumers is small. On a scale from zero to one, where zero corresponds

Author Note

Thanks to Brian Weeks for thoughtful comments on an earlier draft of this manuscript. This material is based in part upon work supported by the National Science Foundation under Grant Numbers IIS-1149599. Any opinions, findings, and conclusions or recommendations expressed in this material are those

of the author(s) and do not necessarily reflect the views of the National Science Foundation.

* Correspondence concerning this article should be addressed to R. Kelly Garrett, School of Communication, Ohio State University, 3016 Derby Hall, 154. N. Oval Mall, Columbus, OH 43210, United States. Contact: garrett.258@osu.edu

to no segregation (e.g., being randomly exposed to news), the overall segregation level was only 0.11 (Flaxman, Goel, & Rao, 2016, p. 308). Further, the authors found that although most Americans rely on only a few online news outlets, the most used outlets tend to attract comparable numbers of conservatives and liberals, suggesting that they represent a more diverse range of views. News outlets that appeal to the political extremes receive relatively little attention. The second study contrasts the ideological segregation that Americans encounter in their online news exposure to what they encounter in a variety of other settings. Analyzing Internet tracking data collected from 12,000 comScore panelists over a 12-month period, this study found that face-to-face interactions tended to be more segregated than online news use (Gentzkow & Shapiro, 2011). In short, the notion that people have constructed highly polarized online news environments, environments in which they never see the other side, is a myth.

Nor are exposure-based echo chambers likely to emerge, as their existence is inconsistent with well documented human information-selection preferences. Selective exposure research dating back to the 1960s has shown that individuals are attracted to attitude-congruent information, but that their response to attitude-discrepant information is more complicated (Frey, 1986; Hart et al., 2009). Importantly, when presented with a mixture of one- and two-sided news stories, individuals do not choose attitude-congruent information at the expense of attitude-discrepant information (Garrett & Stroud, 2014). As long as an individual’s views are represented in a news story, most people are indifferent to the inclusion of other viewpoints. Furthermore, there are instances when the inclusion of other perspectives is preferred (Carnahan, Garrett, & Lynch, 2016; Knobloch-Westerwick & Kleinman, 2011). These patterns have not changed despite the radical transformation of communication technologies (Garrett, Carnahan, & Lynch, 2013).

In a similar vein, Lewandowsky et al. also claim that “most online users are, knowingly or not, put into a filter bubble,” where software systematically shields them from views with which they might disagree (2017, p. 353). Although there is less research on this question, a growing body of evidence suggests that this is a misleading characterization of the influence of what computer scientists often refer to as “recommender systems.” The first evidence comes from the three-month observational study of online news users described above. Not only were overall segregation levels low, but individuals who got their news through technologies commonly characterized as promoting filter bubbles—including social networks and search engines—saw an *increase* in their exposure to news representing the political opposition (Flaxman et al., 2016, p. 316).

A study of over 10 million Facebook users provides additional evidence on this question (Bakshy, Messing, & Adamic, 2015). Researchers at the company compared users’ self-reported ideology to the political orientation of the news that they encountered via the social networking service. The results provide straightforward evidence that Facebook does not systematically screen out all counter-attitudinal exposure. Algorithmic filtering removed a relatively small fraction of the cross-cutting content found in users’ news feeds, reducing it by 5% for conservatives

and 8% for liberals. After filtering, more than one in five political stories to which Facebook users in the study were exposed was cross-cutting.

Although it is obviously possible to tailor automated recommender systems to promote filter bubbles, designers have long recognized the importance of diversity and serendipity (Negroponte, 1996), and have strived to protect it. Over the past decade, scholars interested in system design have proposed and tested dozens of approaches that actively promote diverse information exposure (for reviews, see Bozdag & Hoven, 2015; Garrett & Resnick, 2011). Furthermore, since recommender systems typically “learn” users’ preferences by extrapolating from past behavior, the psychological tendency to tolerate, and sometimes seek, counter-attitudinal information should help preserve diversity in recommended content.

This does not mean that the authors are wrong to assert that the online environment has contributed to the rise of a post-truth mentality, but it does mean that some of the mechanisms they describe are misspecified. Different mechanisms suggest different remedies. If ignorance induced by echo chambers or filter bubbles were the problem, then diversifying exposure would be an obvious solution. But something else is going on. Given the relative rarity of echo chambers, strategies focused on countering them are unlikely to have a substantively important influence on belief accuracy. It is crucial that we understand why these inaccurate beliefs persist so that we can develop strategies that target those causes specifically.

Reconceptualizing the Threat

The terms “echo chamber” and “filter bubble” are sometimes used in a way that is empirically supported. The terms have been used to describe social media practices that exhibit highly segmented content *engagement*—rather than exposure—in the form of “likes,” shares, and comments (Schmidt et al., 2017). For example, someone who “likes” a Facebook post about a conspiracy theory is unlikely to engage substantively with other more scientifically informed posts (Zollo et al., 2017).

Both the antecedents and consequences of exposure differ substantively from those of engagement. For example, an individual who is motivated to consume counter-attitudinal content may be much less likely to endorse that same content with a Facebook like (Stroud, Muddiman, & Scacco, 2013). And reading a news article that challenges one’s beliefs is less likely to be persuasive when it is accompanied by comments (from a like-minded individual) that challenge the articles’ conclusions (Anderson, Brossard, Scheufele, Xenos, & Ladwig, 2014). Given these differences, the use of “echo chambers” to refer to both exposure and engagement is potentially confusing. For this reason, I use the label “engagement echo chambers” when referring to highly segmented interaction with social media content.

Engagement echo chambers can promote falsehoods regardless of the diversity of information exposure. One does not need to avoid contact with all belief-incongruent information to maintain inaccurate beliefs (Kahan, 2015). Belonging to a social network that consistently affirms one view can promote

endorsement of that view, regardless of what other information is encountered. There are a variety of mechanisms by which this occurs. For some, endorsing a falsehood has social utility, for example by allowing the individual to avoid being ostracized by the in-group (Kahan, 2013). Repeated exposure to expressions of support for an idea (regardless of the evidence encountered) can also contribute to a false consensus effect, leading individuals to overestimate public support for their beliefs. (This is a subtle variation on the authors’ description of false consensus resulting from one-sided news exposure.)

These dynamics align nicely with Lewandowsky et al.’s (2017) warning that social media can amplify the effects of narrowly targeted political deception. I argue, however, that concerns about any form of echo chamber are a distraction from a much larger threat. As a society, we should be more concerned about how technologies are used by political strategists, private interests, and foreign powers to manipulate people for political gain than by self-imposed, or technology-induced, political isolation. In other words, we should be focused on finding ways to fight disinformation campaigns, which often pair deceptive messages and strategic extremity.

Using misleading or false claims to sway public opinion dates back to the time of ancient Greece (Mara, 2008), but its profile in the U.S. has grown rapidly in the past two decades. Strategically deployed falsehoods have played an important role in shaping Americans’ attitudes toward a variety of high-profile political issues. The oil industry has worked to promote doubt about climate change science using tactics pioneered by cigarette manufacturers in the 1960s (Jervig, Jennings, Hirsch, & Rust, 2015; Oreskes & Conway, 2010). Widespread misperceptions promoted by the Bush administration about the evidence for WMDs in Iraq and about Iraq’s role in the 9/11 terrorist attacks helped motivate public support for the U.S.-led invasion (Prasad et al., 2009; World Public Opinion, 2006). Deceptive characterization of the Affordable Care Act, including the infamous “death panel” claims have contributed to negative attitudes toward the law (Meirick, 2013; Nyhan, 2010). False claims about political candidates abound, as evidenced by the repeated questioning of President Obama’s birthplace and religion by his political opponents while he was in office (Hartman & Newmark, 2012; YouGov Staff, 2014). In short, deceptive political practices are on the rise, precipitating what has been aptly described as “the demise of ‘fact’ in political discourse” (Jamieson, 2015).

The 2016 U.S. Presidential election is particularly alarming for several reasons. President Trump made an extraordinary number of false or unsubstantiated claims during the campaign (Holan, 2016; Swire, Berinsky, Lewandowsky, & Ecker, 2017). Perhaps more importantly, fact checking had little influence on the behavior of then-candidate Trump, which is a discouraging contrast to prior empirical evidence on the subject (Nyhan & Reifler, 2015). Another source of concern about the 2016 election is evidence that the Russian government sought to use disinformation to influence the outcome of the election (Office of the Director of National Intelligence, 2017; and see Paul & Matthews, 2016). These efforts extended to information

operations conducted on Facebook (Weedon, Nuland, & Stamos, 2017), and were likely responsible, at least in part, for the widely read “fake news” that received so much coverage in the last few months of the campaign season (e.g., Silverman, 2016).

Technology for Fighting Disinformation

Shifting blame for misperceptions from ignorance induced by echo chambers to strategic disinformation campaigns has implications for which types of technologies will be most effective. Focusing on this issue suggests three strategies not considered by Lewandowsky et al. (2017). The first approach is grounded in the observation that although it may be difficult to convince opinion leaders to update their inaccurate beliefs, educating large numbers of individuals who are less influential but more receptive to belief change can also precipitate significant shifts in public opinion (Watts & Dodds, 2007). Many people are only weakly committed to their beliefs, even on highly publicized and extremely contentious issues. For instance, three in ten American say they could easily change their mind about climate change (Leiserowitz et al., 2014). Although any one of these individuals by definition has little influence, shifts toward more accurate beliefs among large numbers of them can be a powerful persuasive force (Watts & Dodds, 2007). Thus, targeting educational campaigns at those most receptive to belief change, and encouraging those individuals to share their knowledge with other members of their social network has significant potential.

Technology may be able to help both by delivering accurate information to those most open to belief change and by helping them to share their accurate beliefs widely. There are efforts afoot to automatically flag claims that fact checkers have determined to be inaccurate (Costine, 2016), and these technologies could provide a foundation for other, potentially more influential capabilities. Rather than just flagging inaccurate content, for example, a social media service could encourage recipients to review the corrections they receive, and share those they find persuasive with the original poster. Involving multiple individuals belonging to the poster’s social network brings contextual awareness to the evaluation process, and it could help reduce the reactance that individuals sometimes experience when presented with automated corrections (see Garrett & Weeks, 2013). Information quality may not be a strong predictor of which information individuals are predisposed to share, but learning that a claim is false can reduce sharing, even on social media. For example, individuals often respond to a well sourced peer correction, such as a link to Snopes, by removing the offending post (Friggeri, Adamic, Eckles, & Cheng, 2014).

Another technology-centric approach focuses on using search engines to reduce exposure to strategic falsehoods in the first place. Anytime we consider using technologies to shape the flow of information on a large scale we must be sensitive to the risk of censorship, but given that we have a legitimate interest in suppressing disinformation campaigns, it is an option worth considering. In other words, it is reasonable to at least ask whether we can find ways to stem the flow of strategic deception

without undermining the free expression of ideas. Search engine users are highly influenced by result rankings (Pan et al., 2007), leading scholars to worry about the possible political influence of the underlying algorithms (Epstein & Robertson, 2015). But such influence could be put to good use. Google is exploring at least two approaches for fighting disinformation. First, it encourages organizations engaged in fact checking to make their results machine readable so that Google can integrate the conclusions into search results (Kosslyn & Yu, 2017). The result is visible as of this writing: search for “Obama Kenyan citizen,” and Google will provide a short statement summarizing the false claim alongside an assessment: “Fact check by Snopes.com: FALSE” (Weise, 2017). Second, Google is developing technology intended to algorithmically estimate source trustworthiness (Luna Dong et al., 2015). These estimates can then be used to adjust rankings, so that unreliable sources have a lower profile (Hodson, 2015).

Introducing claim accuracy and source reliability into search engines constitutes a new form of algorithmic gatekeeper in the digital media landscape, one that is likely to help stem the flow of disinformation. By restoring barriers to reaching a mass audience that were greatly reduced by the explosive growth of online news, these software tools could provide a social good, much as respected newspaper editors and television news anchors did in an earlier era. Given the difficulty of unseating congenial falsehoods—especially when they are seen by large numbers of people predisposed to believe them—reducing exposure to socially harmful falsehood is valuable.

The ethical challenges are undeniable. It may be necessary to curb the flow of false claims if we want truth to prevail, but how do we achieve this without constraining dissent or legitimate debate over what is true? We need to balance attempts to reduce exposure to socially harmful content with protection of minority opinion. This task has precedent, both legal (e.g., Sunstein, 2009) and philosophical (e.g., Möllering, 2009). Still, more work is needed. It will require conceptual work to articulate the boundary between constraining the spread of misinformation and censorship, as well as empirical research to ensure that automated systems can effectively distinguish between falsehood and dissent.

The first two strategies, one focusing more on social media and the other more on search engines, are closely interrelated. They concern ways that technology can influence the flow of information, promoting the dissemination of accurate claims while constraining the flow of falsehoods. A less direct approach builds on the fact that disinformation campaigns often use emotional extremity for strategic effect. Messages that convey anger, outrage, and distrust can have a powerful influence on how people respond to their information environment. Anger increases individuals’ susceptibility to in-group falsehoods (Weeks, 2015), as does animosity and distrust toward the out-group (Lodge & Taber, 2013). The vitriol found in online comments can undermine the credibility of the news source (Thorson, Vraga, & Ekdale, 2010). One potential contribution,

then, is to find ways to slow or reverse a thirty-plus year trend toward increasing affective polarization, the hostility that partisans feel toward those on the other side of the aisle (Iyengar, Sood, & Lelkes, 2012).

There is precedent for proposing sociotechnical solutions to similar problems (e.g., the decline in social capital; Resnick, 2004), and there are many ways that we might tackle this one. Rather than attempting to list all possible strategies, I highlight a few areas of relevant on-going research and describe a novel strategy tied to people’s use of social media today. Scholars interested in online deliberation have created technologies designed to help individuals discuss controversial topics in productive ways (for a review, see Friess & Eilders, 2015). These tools have been shown to improve outcomes and to make people feel more positively about their engagement with difficult topics. However, they do nothing to prevent strategic efforts to promote anger and distrust that are so common today (Berry & Sobieraj, 2013; Jamieson & Cappella, 2008). One approach worth pursuing is to promote more deliberation in the spaces where online discussion is already taking place. For example, having journalists contribute to their news organization’s Facebook page can promote civility (Stroud, Scacco, Muddiman, & Curry, 2015). It might also be possible to adapt some of the features of existing online deliberation systems for use in common conversation spaces, including social media. For example, software that prompts users to respond to others comments by first explaining in their own words what they believe was said can promote better (online) conversation (Kriplean, Toomim, Morgan, Boring, & Ko, 2012).

There is a variation on this strategy that has both greater risks and greater potential rewards. Given the volume of content shared online today, social media systems already make numerous choices about what to show, and in what order. These decisions are typically informed by factors such as popularity within an individual’s social network, and the user’s previous engagement with similar content. Suppose that content filters also considered the emotional tenor of messages when deciding what to present to users. Emotions can spread across social networks (Kramer, Guillory, & Hancock, 2014), and angry or hateful messages could be an important conduit. By reducing the profile of the most noxious contributions, it may be possible to slow the spread of politically motivated vitriol. This would in turn lessen the risk that individuals will succumb to deception, but like efforts to automatically limit the flow of falsehoods directly, such a strategy poses important ethical questions. The idea that one must respect conversational conventions (e.g., being polite) in order to participate in discussion is not new. To the contrary, the anything-goes communication style so often found online is a relatively recent phenomenon. Using machines to police rules about what constitutes permissible speech could, however, have both good and bad consequences. Social conventions are sometimes used to silence dissent (Mansbridge, 1999), while the absence of conventions can be used to promote hatred and violence. Pursuing this course of action requires caution.

Conclusion

Designing technological interventions intended to promote a more knowledgeable, less easily deceived citizenry demands that we attend carefully to the social and psychological factors that shape the flow and acceptance of misinformation. Lewandowsky et al. (2017) have done a great service by providing a broad overview of relevant phenomena. Although their descriptions are careful, they are also necessarily brief. As a consequence, there is some risk that readers will misinterpret references to “echo chambers” and “filter bubbles.” In this response, I have sought to show that online news users are not insulated from information with which they disagree, making this an unlikely source of misperception. More important than exposure-based echo chambers are strategic efforts to use misinformation and outrage to promote political objectives.

Social resources for fighting disinformation are finite, and we cannot afford to be distracted by trying to solve problems that occur only rarely. Instead, we must focus our attention on efforts to use social media and search engines in the service of disinformation campaigns. Following Lewandowsky et al.'s (2017) example, I briefly suggest three additional strategies that may help: helping users correct their peers, accounting for accuracy in search engine results, and using technology to promote civility and to slow the spread of hostile or polarizing content.

Conflict of Interest Statement

The authors declare no conflict of interest.

Keywords: Misinformation, Disinformation, Echo chamber, Filter bubble, Information, Communication, Technology, Recommender system

References

- Allport, G. W., & Postman, L. (1965 [1947]). *The psychology of rumor*. New York: Russell & Russell.
- Anderson, A. A., Brossard, D., Scheufele, D. A., Xenos, M. A., & Ladwig, P. (2014). The “nasty effect:” Online incivility and risk perceptions of emerging technologies. *Journal of Computer-Mediated Communication*, 19(3), 373–387. <http://dx.doi.org/10.1111/jcc4.12009>
- Bakshy, E., Messing, S., & Adamic, L. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239), 1130–1132. <http://dx.doi.org/10.1126/science.aaa1160>
- Berry, J. M., & Sobieraj, S. (2013). *The outrage industry: Political opinion media and the new incivility*. New York: Oxford University Press.
- Bozdag, E., & Hoven, J. (2015). Breaking the filter bubble: Democracy and design. *Ethics and Information Technology*, 1–17. <http://dx.doi.org/10.1007/s10676-015-9380-y>
- Carnahan, D., Garrett, R. K., & Lynch, E. (2016). Candidate vulnerability and exposure to counter-attitudinal information: Evidence from two U.S. presidential elections. *Human Communication Research*, 42(4), 577–598. <http://dx.doi.org/10.1111/hcre.12088>
- Costine, J. (2016). *Facebook now flags and down-ranks fake news with help from outside fact checkers*. Retrieved from <https://techcrunch.com/2016/12/15/facebook-now-flags-and-down-ranks-fake-news-with-help-from-outside-fact-checkers/>
- Epstein, R., & Robertson, R. E. (2015). The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *Proceedings of the National Academy of Sciences of the United States of America*, 112(33), E4512–E4521. <http://dx.doi.org/10.1073/pnas.1419828112>
- Flaxman, S., Goel, S., & Rao, J. M. (2016). Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly*, 80(S1), 298–320. <http://dx.doi.org/10.1093/poq/nfw006>
- Frey, D. (1986). Recent research on selective exposure to information. *Advances in Experimental Social Psychology*, 19, 41–80.
- Friess, D., & Eilders, C. (2015). A systematic review of online deliberation research. *Policy & Internet*, 7(3), 319–339. <http://dx.doi.org/10.1002/poi3.95>
- Frigerri, A., Adamic, L., Eckles, D., & Cheng, J. (2014). Rumor cascades. In *Paper presented at the eighth international AAAI conference on weblogs and social media, North America* <http://www.aaai.org/ocs/index.php/ICWSM/ICWSM14/paper/view/8122>
- Garrett, R. K., Carnahan, D., & Lynch, E. K. (2013). A turn toward avoidance? Selective exposure to online political information, 2004–2008. *Political Behavior*, 35(1), 113–134. <http://dx.doi.org/10.1007/s11109-011-9185-6>
- Garrett, R. K., & Resnick, P. (2011). Resisting political fragmentation on the Internet. *Daedalus*, 140(4), 108–120. http://dx.doi.org/10.1162/DAED_a.00118
- Garrett, R. K., & Stroud, N. J. (2014). Partisan paths to exposure diversity: Differences in pro- and counterattitudinal news consumption. *Journal of Communication*, 64(4), 680–701. <http://dx.doi.org/10.1111/jcom.12105>
- Garrett, R. K., & Weeks, B. E. (2013, February). The promise and peril of real-time corrections to political misperceptions. In *Paper presented at the CSCW '13 proceedings of the 2013 conference on computer supported cooperative work*.
- Genzkow, M., & Shapiro, J. M. (2011). Ideological segregation online and offline. *Quarterly Journal of Economics*, 126(4), 1799–1839. <http://dx.doi.org/10.1093/qje/qjr044>
- Hart, W., Albarracín, D., Eagly, A. H., Brechan, I., Lindberg, M. J., & Merrill, L. (2009). Feeling validated versus being correct: A meta-analysis of selective exposure to information. *Psychological Bulletin*, 135(4), 555–588. <http://dx.doi.org/10.1037/a0015701>
- Hartman, T. K., & Newmark, A. J. (2012). Motivated reasoning, political sophistication, and associations between president Obama and Islam. *PS: Political Science & Politics*, 45(3), 449–455. <http://dx.doi.org/10.1017/S1049096512000327>
- Hodson, H. (2015). Nothing but the truth. *New Scientist* 2015., NEWS; No. 1234, February 28.
- Holan, A. D. (2016). *2016 lie of the year: Fake news.*. Retrieved from <http://www.politifact.com/truth-o-meter/article/2016/dec/13/2016-lie-year-fake-news/>
- Iyengar, S., Sood, G., & Lelkes, Y. (2012). Affect, not ideology. *Public Opinion Quarterly*, 76(3), 405–431. <http://dx.doi.org/10.1093/poq/nfs038>
- Jamieson, K. H. (2015). Implications of the demise of “fact” in political discourse. *Proceedings of the American Philosophical Society*, 159(1), 66–84.
- Jamieson, K. H., & Cappella, J. N. (2008). *Echo chamber: Rush Limbaugh and the conservative media establishment*. Oxford: Oxford University Press.

- Jerving, S., Jennings, K., Hirsch, M. M., & Rust, S. (2015). *What Exxon knew about the Earth's melting arctic*. Los Angeles Times. Retrieved from. <http://graphics.latimes.com/exxon-arctic/>
- Kahan, D. M. (2013). Ideology, motivated reasoning, and cognitive reflection: An experimental study. *Judgment and Decision Making*, 8, 407–424.
- Kahan, D. M. (2015). Climate science communication and the measurement problem. *Advances in Political Psychology*, 36, 1–43. <http://dx.doi.org/10.1111/pops.12244>
- Knapp, R. H. (1944). A psychology of rumor. *Public Opinion Quarterly*, 8(1), 22–37.
- Knobloch-Westerwick, S., & Kleinman, S. (2011). Pre-election selective exposure: Confirmation bias versus informational utility. *Communication Research*, 39(2), 170–193. <http://dx.doi.org/10.1177/0093650211400597>
- Kosslyn, J., & Yu, C. (2017). *Fact check now available in Google Search and news around the world*. Retrieved from. <https://www.blog.google/products/search/fact-check-now-available-google-search-and-news-around-world/>
- Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences of the United States of America*, 111(24), 8788–8790. <http://dx.doi.org/10.1073/pnas.1320040111>
- Kriplean, T., Toomim, M., Morgan, J., Borning, A., & Ko, A. (2012). Is this what you meant?: Promoting listening on the web with reflect. In *Paper presented at the proceedings of the SIGCHI conference on human factors in computing systems*.
- Leiserowitz, A., Maibach, E., Roser-Renouf, C., Feinberg, G., Rosenthal, S., & Marlon, J. (2014). *Climate change in the American mind: American's global warming beliefs and attitudes in November 2013*. New Haven, NJ: Yale Project on Climate Change Communication. Retrieved from. <http://environment.yale.edu/climate-communication/files/Climate-Beliefs-November-2013.pdf>
- Lewandowsky, S., Ecker, U., & Cook, J. (2017). Beyond misinformation: Understanding and coping with the post-truth era. *Journal of Applied Research in Memory and Cognition*.
- Lodge, M., & Taber, C. S. (2013). *The rationalizing voter*. New York, NY: Cambridge University Press.
- Luna Dong, X., Gabrilovich, E., Murphy, K., Dang, V., Horn, W., Lugaresi, C., et al. (2015). *Knowledge-based trust: Estimating the trustworthiness of web sources*. Retrieved from. <http://arxiv.org/abs/1502.03519>
- Mansbridge, J. (1999). Everyday talk in the deliberative system. In S. Macedo (Ed.), *Deliberative politics: Essays on democracy and disagreement* (pp. 211–239). New York: Oxford University Press.
- Mara, G. M. (2008). *The civic conversations of Thucydides and Plato: Classical political philosophy and the limits of democracy*. Albany: SUNY Press.
- Meirick, P. C. (2013). Motivated misperception? Party, education, partisan news, and belief in “death panels”. *Journalism & Mass Communication Quarterly*, <http://dx.doi.org/10.1177/1077699012468696>
- Möllering, G. (2009). Leaps and lapses of faith: Exploring the relationship between trust and deception. In B. Harrington (Ed.), *Deception: From ancient empires to Internet dating* (pp. 137–153). Stanford, CA: Stanford University Press.
- Negroponte, N. (1996). *Being digital*. New York: Vintage Books.
- Nyhan, B. (2010). Why the “death panel” myth wouldn’t die: Misinformation in the health care reform debate. *The Forum*, 8(1) <http://dx.doi.org/10.2202/1540-8884.1354>. Article 5
- Nyhan, B., & Reifler, J. (2015). The effect of fact-checking on elites: A field experiment on U.S. state legislators. *American Journal of Political Science*, 59(3), 628–640. <http://dx.doi.org/10.1111/ajps.12162>
- Office of the Director of National Intelligence. (2017). *Assessing Russian activities and intentions in recent US Elections Intelligence Community Assessment*. Retrieved from. [https://web-beta.archive.org/web/20170421222356/; https://www.dni.gov/files/documents/ICA_2017_01.pdf](https://web-beta.archive.org/web/20170421222356/;https://www.dni.gov/files/documents/ICA_2017_01.pdf)
- Oreskes, N., & Conway, E. M. (2010). *Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming*. New York: Bloomsbury Publishing.
- Pan, B., Hembrooke, H., Joachims, T., Lorigo, L., Gay, G., & Granka, L. (2007). In Google we trust: Users’ decisions on rank, position, and relevance. *Journal of Computer-Mediated Communication*, 12(3), 801–823. <http://dx.doi.org/10.1111/j.1083-6101.2007.00351.x>
- Paul, C., & Matthews, M. (2016). *The Russian “firehose of falsehood” propaganda model: Why it might work and options to counter it perspectives*. pp. 16. Rand Corporation. Retrieved from <http://www.rand.org/pubs/perspectives/PE198.html>
- Prasad, M., Perrin, A. J., Bezila, K., Hoffman, S. G., Kindleberger, K., Manturuk, K., et al. (2009). “There must be a reason”: Osama, Saddam, and inferred justification. *Sociological Inquiry*, 79(2), 142–162. <http://dx.doi.org/10.1111/j.1475-682X.2009.00280.x>
- Resnick, P. (2004). Impersonal sociotechnical capital, ICTs, and collective action among strangers. In W. Dutton, B. Kahin, R. O’Callaghan, & A. Wyckoff (Eds.), *Transforming enterprise*. MIT Press.
- Schmidt, A. L., Zollo, F., Del Vicario, M., Bessi, A., Scala, A., Caldarelli, G., et al. (2017). Anatomy of news consumption on Facebook. *Proceedings of the National Academy of Sciences of the United States of America*, 114(12), 3035–3039. <http://dx.doi.org/10.1073/pnas.1617052114>
- Shibutani, T. (1966). *Improvised news: A sociological study of rumor*. Indianapolis: Bobbs-Merrill.
- Silverman, C. (2016). *This analysis shows how viral fake election news stories outperformed real news on Facebook*. Retrieved from. <https://www.buzzfeed.com/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook>
- Stroud, N. J., Muddiman, A., & Scacco, J. M. (2013). *Engaging news project: Social media buttons in comment sections*. Austin, TX: Annette Strauss Institute for Civic Life. Retrieved from. http://engagingnewsproject.org/enp_prod/wp-content/uploads/2014/04/ENP_Buttons_Report.pdf
- Stroud, N. J., Scacco, J. M., Muddiman, A., & Curry, A. L. (2015). Changing deliberative norms on news organizations’ Facebook sites. *Journal of Computer-Mediated Communication*, 20(2), 188–203. <http://dx.doi.org/10.1111/jcc4.12104>
- Sunstein, C. R. (2009). *On rumors: How falsehood spread, why we believe them, what can be done*. New York, NY: Farrar, Straus and Giroux.
- Swire, B., Berinsky, A. J., Lewandowsky, S., & Ecker, U. K. H. (2017). Processing political misinformation: Comprehending the Trump phenomenon. *Royal Society Open Science*, 4(3) <http://dx.doi.org/10.1098/rsos.160802>
- Thorson, K., Vraga, E., & Ekdale, B. (2010). Credibility in context: How uncivil online commentary affects news credibility. *Mass Communication and Society*, 13(3), 289–313. <http://dx.doi.org/10.1080/15205430903225571>
- Watts, D. J., & Dodds, P. S. (2007). Influentials, networks, and public opinion formation. *Journal of Consumer Research*, 34(4), 441–458. <http://dx.doi.org/10.1086/518527>

- Weedon, J., Nuland, W., & Stamos, A. (2017). *Information operations and Facebook*. Facebook, Inc. Retrieved from. <https://fbnewsroomus.files.wordpress.com/2017/04/facebook-and-information-operations-v1.pdf>
- Weeks, B. E. (2015). Emotions, partisanship, and misperceptions: How anger and anxiety moderate the effect of partisan bias on susceptibility to political misinformation. *Journal of Communication*, 65(4), 699–719. <http://dx.doi.org/10.1111/jcom.12164>
- Weise, E. (2017). *We tried Google's new fact-check filter on the Internet's favorite hoaxes*. USA Today. Retrieved from. <https://www.usatoday.com/story/tech/news/2017/04/10/google-fact-check-snopes-politifact-factcheck/100263464/>
- World Public Opinion. (2006 August). *Percentage of Americans believing Iraq had WMD rises..* Retrieved from. <http://worldpublicopinion.net/percentage-of-americans-believing-iraq-had-wmd-rises/>
- YouGov Staff. (2014). *15% 'know for sure' Obama was not born in the U.S. Economist/YouGov poll*. Retrieved from. <https://today.yougov.com/news/2014/02/12/know-for-sure/>
- Zollo, F., Bessi, A., Del Vicario, M., Scala, A., Caldarelli, G., Shekhtman, L., et al. (2017). Debunking in a world of tribes. *PLoS ONE*, 12(7), e0181821. <http://dx.doi.org/10.1371/journal.pone.0181821>

Received 13 September 2017;
received in revised form 27 September 2017;
accepted 28 September 2017