The cognitive science of fake news.

Neil Levy^{1*} and Robert M. Ross¹

¹Department of Philosophy, Macquarie University.

*Corresponding author: neil.levy@mq.edu.au

24th August 2020

Forthcoming book chapter to appear in:

Hannon, M. & de Ridder, J. (2021). *Routledge Handbook of Political Epistemology*.

Routledge.

This manuscript has not yet been copy-edited and the published book chapter might have minor differences.

Abstract

In this chapter, we provide a necessarily brief and partial survey of recent work in the cognitive sciences directly on or closely related to the psychology of fake news, in particular fake news in the political domain. We focus on whether and why people believe fake news. While we argue that it is likely that a large proportion of people who purport to believe fake news really do, we provide evidence that this proportion might be significantly smaller than is usually thought (and smaller than is suggested by surveys). Assertion of belief is inflated, we suggest, by insincere report, whether to express support for one side of political debate or simply for fun. It is also inflated by the use of motivated inference of one sort or another, which lead respondents to report believing things about which they had no opinion prior to being probed. We then turn to rival accounts that aim to explain why people believe in fake news when they do. While partisan explanations, turning on motivated reasoning, are probably best known, we show they face serious challenges from accounts that explain belief by reference to analytic thinking.

Fake news is, roughly, the set of reports of events of public interest (“news”) that purport to be or which mimic reliable news sources but which intend to deceive or are indifferent to truth.¹ Fake news is very widely disseminated: for example, during the 2016 US presidential election, the most popular fake stories were shared more widely than the most popular genuine new stories (Vosoughi et al., 2018). While the extent of its consumption and its influence over behaviour is often exaggerated (Allen et al., 2020; Guess et al., 2020c, 2020b; Mercier, 2020), it is plausible that it is big enough to be decisive when matters are finely balanced (in close elections, for instance). Some people credit, or blame, fake news with swinging the Brexit referendum (Grice, 2017) and the 2016 election of Donald Trump (Gunther et al., 2018).

While there are a range of purely conceptual questions raised by fake news, fully understanding and addressing its rise, reach and influence requires grappling with a range of empirical findings. The cognitive science of fake news specifically is in its infancy, but there is a rich literature on related topics – on misinformation and its effects, conspiratorial ideation, and so on – on which to draw. The questions asked by cognitive scientists are often continuous with those raised by philosophers, and also provoke further philosophical reflection. In this chapter, we will survey and critically assess some of this rich literature. We note at the outset that our discussion is heavily skewed to the US context, because most of the available research concerns the US. That is an important limitation of this discussion, since the US may have features that entail that fake news has different effects in that context than elsewhere.

There are at least five questions concerning the spread and influence of fake news that might be addressed by, or with the help of, cognitive science: (1) why do people share fake news? (2) to what extent do they believe it? (3) on the assumption they do believe it, what explains their belief? (4) what influence does it have over their behaviour? (5) to the extent to which its influence is negative, either epistemically or behaviourally, how can we reduce these impacts? In the interest of space, we will concentrate on (2), (3) and (5); even so restricted, our discussion will necessarily be highly selective.

Do people believe fake news?

A surprisingly high number of people express agreement with some of the more widely disseminated fake news items. For instance, one survey found that approximately a third of Americans reported believing the “Birther” theory that Barack Obama was not born in the United States (Uscinski and Parent, 2014). However, while assertion is usually a very reliable guide to belief (and may be the most powerful cue to belief attribution; Rose et al., 2014), it is not infallible. There is evidence of a belief/behaviour mismatch on politically charged topics, which calls into question the sincerity of people’s belief reports. McGrath (2017) found little evidence that political partisans changed their economic behaviour in line with their assertions about the effects of preferred and non-preferred candidates on the economy. However, other studies have found evidence that people do tend to behave in ways that conform to their professed beliefs. For instance, Lerman et al. (2017) found that Republicans do not merely say that “Obamacare” is undesirable; they also enrol in it at lower rates than those who express more favourable attitudes. Because these data are difficult to interpret and remain sparse, we will not say more about them here. Note, though, that our discussion

of survey responses has direct relevance to claims about how we ought to interpret behavioural evidence. That is as it should be: assertion is, after all, behaviour.

Political attitudes and beliefs, including beliefs in misinformation, have primarily been probed through survey instruments. Surveys have long documented large partisan gaps in beliefs, with each side perceiving the world in a way that seems to conform to their normative views (Jerit and Barabas, 2012; note, however, that Roush and Sood, 2020 have recently provided evidence that partisan knowledge gaps might be much smaller than previously thought). Thus, Republicans and Democrats report diverging beliefs about factual claims (e.g. the unemployment rate). However, there is an alternative explanation of these reports: perhaps they are not sincere reports of beliefs at all, but more akin to partisan cheerleading (Bullock and Lenz, 2019). That is, rather than report their sincere beliefs, people may seek to *express* their allegiance to a party, a policy or a person. Someone might assert “Obama is a secret Muslim” not because they believe he is a Muslim, but rather to express their dislike of him. Since these responses have an expressive function, we will adopt the terminology used by other researchers (Berinsky, 2018; Bullock et al., 2015) and call this kind of behaviour *expressive responding*.

There is persuasive evidence that people engage in expressive responding (Hannon, 2019). The Trump administration notoriously claimed that the crowd at the incoming president’s inauguration was the largest ever, claims that flew in the face of the photographic evidence that clearly showed that Obama’s inauguration in 2009 was much larger. On the days immediately following the controversy, Schaffner and Luks (2018) presented participants with photos of the two events and asked which depicted a larger crowd, knowing that many

participants would recognize them as photos of the Trump and Obama inaugurations. A very small proportion of non-voters and Clinton voters identified the photo of the Trump inauguration as depicting a larger crowd (3% and 2% respectively). This contrasts with the response of Trump voters: 15% of them identified the photo of his inauguration as depicting a larger crowd. This strongly suggests that some individuals are willing to report a belief they do not hold in order to express support for their preferred party or candidate. Schaffner and Luks suggest that this figure may represent a lower bound for expressive responding: some Trump supporters were probably unaware of the controversy, and therefore did not see the task as presenting them with an opportunity for an expressive response.

Evidence from other studies, using different methodologies to attempt to estimate the prevalence of expressive responding, seems to indicate this study is an outlier in terms of the magnitude of the effect (perhaps because conditions are rarely so ideal for expressive responding). Prior et al. (2015) found that the possibility of small monetary rewards for correct responses halved partisan bias (from 12% to 6%). Bullock et al. (2015) report similar results, and an apparent dose dependence of reduction: the larger the incentive, the bigger the reduction in bias; and a combination of treatments eliminated partisan differences altogether. In contrast, Berinsky (2017) found little or no evidence of expressive responding, despite offering an incentive (albeit one of a different type: a reduction in time spent on the survey); and Peterson and Iyengar (2020) found that a partisan gap about two thirds of the unincentivized gap persisted with financial incentives. Taken together, and despite some failures to narrow the partisan gap via the provision of incentives, the evidence strongly suggests that a substantial number of respondents express attitudes, rather than report beliefs, in surveys (see Bullock and Lenz (2019) for a comprehensive review). It is important

to note, however, that only Berinsky probed beliefs of the kind that tend to feature in fake news (e.g. 9/11 “truther” claims) so we should be wary of generalizing from evidence of expressive responding in other contexts to the contexts of fake news more specifically.

However, there are reasons to suspect that studies that utilize incentives to measure the prevalence of expressive responding underestimate their extent. Most obviously, people may misrepresent their true beliefs if they know they are controversial, in order to secure the monetary reward. In addition, incentives for accuracy might have perverse effects: someone who wants to express strong support for a politician or a party may do so by spurning the opportunity to receive monetary compensation for reporting their true belief. We might therefore worry that many partisan participants might be difficult to shift by modest monetary reward. Moreover, if participants count support for a person, policy or stance as a sacred value, they may reject the opportunity for financial reward for accuracy: sacred values are usually held to be incommensurable with and tainted by financial reward (Tetlock, 2003). Berinsky’s study might fail to demonstrate expressive responding for one or both these reasons.

In addition, incentivization (and polling more generally) may have the effect of *producing* the very beliefs that are reported. As well as expressive responding, participants may use partisan heuristics (rules of thumb sensitive to cues of party affiliation or of identities and activities valued by or associated with their side of politics), biased sampling methods (e.g. searching memory – or the internet – for evidence that bears on a question, but in a way that is more sensitive to supporting than disconfirming evidence) or other motivated ways of drawing inferences to generate a response, in the absence of a prior belief and without being

confident in the truth of the proposition they assert. While this isn't expressive responding – participants do not report a belief they do not hold, in order to express an attitude – nor is it the veridical report of a *pre-existing* belief (of course, tricky issues to do with dispositional beliefs arise at this point, but we think it is reasonable to restrict the latter to beliefs that are entailed by agents' representations, and thus to exclude beliefs like this, generated by inference from prompts).

Political scientists have long recognized that a substantial proportion of respondents to surveys on political issues construct their responses on the spot (Zaller and Feldman, 1992). Thus, surveys of public opinion play a role in producing the responses they aim to probe. The person who reports believing *that Obama is a secret Muslim* or *that Hilary Clinton gave uranium to Russia in exchange for donations* may not have believed these things prior to being asked. Rather, they may have engaged in biased memory search or biased inference procedures, or applied heuristics, to construct the belief on the spot (while this route to response is less likely to be employed when a claim is widely circulated – agents can be expected to have attitudes on many familiar fake news stories – their current attitude might reflect an earlier reliance on such routes to response). Some support for this view is provided by Bullock et al. (2015) who found that when a “don't know” response was incentivized (at a lower level than for correct responses), use of this option increased dramatically. Similarly, Clifford et al. (2020) found that offering participants realistic alternative options decreased endorsement of conspiratorial explanations markedly. While beliefs constructed for the purposes of response may sometimes persist, surveys that fail to discourage the construction of such beliefs overreport their prevalence in the population.²

Between expressive responding and overreporting as a consequence of belief construction, the extent to which people believe fake news is probably exaggerated by surveys. Nevertheless, there can be little doubt that substantial numbers of people do come to give sufficient credence to fake news for it to affect their behaviour. For instance, in 2017 a false story that the founder of Ethereum had died in a car accident caused the market value of the company to drop by \$4 billion (Dunning, 2019). More disturbingly, a number of Sandy Hook “truthers” – conspiracy theorists who allege that the mass shooting was a false flag operation – have escalated their harassment of parents who lost children at the elementary school beyond online trolling to confronting and threatening them in person (Robles, 2016), and the survivors and relatives of survivors of other mass shootings have also been victims of this kind of activity (Raphelson, 2018). This pattern of behaviour strongly suggests some people come to be convinced of the veracity of some fake news.³ Given that some people come to believe fake news, often in the face of apparent implausibility, and some of them go on to act on these beliefs in ways that are significant and sometimes even harmful to themselves or to loved ones, it is important to understand the mechanisms that explain their acceptance. We turn to this issue next.

Why do people believe fake news?

A natural hypothesis for why people believe fake news, especially when it is prima facie implausible, cites cognitive limitations or ignorance (or, of course, both). However, this deficit model for belief in the implausible or the irrational has come under challenge from a number of researchers, most influentially Dan Kahan. Kahan focuses specifically not on fake news, but on questions like climate change and evolution. These are topics on which there is a large

partisan divide, but one side is at odds with the scientific consensus. They therefore provide an opportunity for probing the processes whereby people come to reject a view that is strongly supported by evidence. On these topics, Kahan (2015) find evidence that a deficit model apparently struggles to explain: while liberals who score higher on the Ordinary Science Intelligence scale (which measures both basic science literacy as well as thinking dispositions) are more likely to accept the science of climate change and the truth of evolution, conservatives who score higher are *less* likely to accept both than conservatives who score less well. Kahan argues that the Ordinary Science Intelligence scale reliably measures cognitive capacity: what matters is not just what capacities agents have, but also how they are deployed. He argues that our values – specifically, our sense of belonging to a particular culture – biases our cognition, such that we reject information sources when they conflict with the beliefs central to that identity (Kahan, 2016). Individuals with greater cognitive capacities also have a greater capacity to explain away inconvenient data. It is easy to see how the deployment of what Kahan terms *identity protective cognition* (Kahan, 2017) might explain acceptance of fake news (or rejection of genuine news).

Kahan's hypothesis appears to be supported by evidence from elsewhere in cognitive science. For instance, there's evidence that better education interacts with political orientation, such that (for example) it predicts higher levels of belief that Barack Obama is a secret Muslim among better educated Republicans (Lewandowsky et al., 2012) and that motivated cognition may enable those with greater capacity or knowledge more effectively to dismiss evidence they find uncongenial (Taber and Lodge, 2006; see Rust and Schwitzgebel, 2009; Schwitzgebel, 2009 for a similar hypothesis in the context of explaining the behaviour of

professional philosophers). Recently, however, Kahan's hypothesis has faced a challenge from theorists who have offered a new version of something like a deficit hypothesis.

Pennycook and Rand (2019a) argue that if susceptibility to fake news were explained by motivated cognition, we ought to see a positive association between belief in politically congenial fake news headlines and cognitive capacities (measured using the Cognitive Reflection Test; CRT). But they found that higher CRT scores correlate with correctly rating fake news as less accurate and genuine news stories as more accurate, regardless of its fit with political leanings. In further work, the same group of researchers provide evidence that higher CRT scores are associated with updating of beliefs about factual political statements closer to Bayesian norms (Tappin et al., 2020a). The focus on belief *updating* is an important innovation because a shortcoming of other study designs is that they do not disentangle political identities from pre-treatment information exposure and issue-specific prior beliefs (Tappin et al., 2020b)

A potential worry for this rival to Kahan's view can be motivated by other work by the very same group of researchers. Their data on voters in the 2016 US presidential election shows an association not only between lower CRT scores and higher support for Trump, but also that lower CRT scores are correlated with lower levels of engagement with politics (Pennycook and Rand, 2019b). Given this association between lower levels of engagement and lower CRT scores, one possible explanation for a decreased capacity to identify fake news (assuming, that is, that those with lower CRT scores genuinely lack such a capacity) is that this subgroup lacked the background knowledge required for assessing the plausibility of the headlines

presented (few were so obviously fake that those without the requisite background could easily dismiss them).

The relative influence of motivated reasoning versus bad reasoning and ignorance on susceptibility to fake news therefore remains an open question for the moment. While there is much more to be said regarding the mechanisms whereby people come to accept fake news, we suggest that the relevant perspectives might best be covered by considering them under the third heading we will discuss next.

How might belief in (and spread of) fake news be prevented or reduced?

As well as looking to cognitive science for explanation of the mechanisms that underlie acceptance of fake news, we might hope to mine cognitive science for strategies to inoculate media consumers against misinformation, or to correct beliefs (or perhaps to reduce reliance on its claims in decision-making).

Obvious candidate interventions for combating fake news are explicit corrections of information reported in fake news and warning labels attached to fake news indicating that the information is not true. There is a large literature examining corrections of misinformation. Meta-analyses of this literature show that corrections do have an effect, but beliefs often continue to be influenced by the misinformation (Chan et al., 2017; Walter et al., 2019; Walter and Tukachinsky, 2020). Given the frequency of belief perseverance, an important focus of research is the properties of corrections that are most effective. Some research suggests that more effective corrections do not simply inform the consumer that

some information is incorrect, but also provide new information that plays the same role as the old (Levy, 2017). For example, one study found that people continue to rely on rebutted information when they have no alternative explanation of an event, but reliance is greatly reduced (though not eliminated) if an explanation is provided (Lewandowsky et al., 2012).

Recent studies have found that a corrective article (Porter et al., 2018) and a warning label (Pennycook et al., 2020a) can reduce belief in fake news. Nonetheless, concerns have been raised about the possibility that corrections and warnings might have negative consequences. At least three classes of negative outcomes have been discussed. First, corrections might elicit a “backfire effect”, with people becoming more committed to a claim following presentation of strong evidence against it. For example, an influential study found that a correction of George W. Bush’s false statement about there being weapons of mass destruction in Iraq before the American invasion increased belief in the false claim among American conservatives (Nyhan and Reifler, 2010). However, follow-up research failed to replicate the backfire effect (Clayton et al., 2019; Porter and Wood, 2019; Wood and Porter, 2019) and a broader reading of the literature suggests that if this phenomena is real, it is not as common as was initially feared (Swire-Thompson et al., 2020). It should be noted however, that Merkley (2020) have recently provided evidence of a backfire effect among those high in anti-intellectualism. A second concern about corrective messages is that they might create an “implied truth effect” where fake news headlines that fail to get tagged as inaccurate are considered validated and more accurate (Pennycook et al., 2020a). This could create significant challenges for factchecking because it isn’t practical for professional fact checkers to check all fake news that is produced. Happily, this research also provides evidence for a solution. The authors also find that attaching verifications to a subset of true headlines

eliminates, and might even reverse, the implied truth effect for untagged headlines. Third, invalid corrections might create a “tainted truth effect” whereby informative news that is wrongly labelled as inaccurate can result in a reduction in the credibility of informative news (Freeze et al., 2020). This could be exploited by bad actors by intentionally creating invalid corrections.

As discussed earlier, individuals high in analytic cognitive ability appear to be better at identifying fake news (Pennycook & Rand, 2019b). A hypothesis that follows naturally from this is that encouraging people to deliberate could reduce belief in fake news. A recent study provides evidence for this (Bago et al., 2020). Participants were presented with a series of partisan news headlines and were asked to make intuitive judgments about the accuracy of the headlines under time pressure and while completing a task designed to tax working memory. They were then given an opportunity to rethink their responses with no constraints, thereby permitting more deliberation and correction of intuitive mistakes. Participants believed fewer false headlines (but not true headlines) when making their final responses. Moreover, other studies from this research group provide evidence that bringing an accuracy motive to the front of people’s minds (by asking them to rate the accuracy of a non-political headline) reduces intention to share fake news on social media (Pennycook et al., 2020b, 2019) and actual sharing of misinformation on Twitter (Pennycook et al., 2019). Together, these studies suggest that interventions that facilitate deliberation on social media platforms could reduce the spread of fake news, such as bringing an accuracy motive to mind by periodically asking users if a randomly selected headline is accurate.

Another potential intervention for reducing belief in fake news is to equip people with skills to identify it. Lutzke et al. (2019) provide experimental evidence that asking oneself four questions can significantly reduce trusting, liking, and intending to share of fake news: (1) *Do I recognize the news organization that posted the story?*; (2) *Does the information in the post seem believable?*; (3) *Is the post written in a style that I expect from a professional news organization?*; and (4) *Is the post politically motivated?* Another study found that digital media literacy intervention increased discernment between mainstream and fake news in both the USA and India (Guess et al., 2020a). However, other research failed to find that making publisher information visible had any influence on the perception of accuracy of unreliable sources or in reducing intention to share fake news (Dias et al., 2020).

Another approach uses “inoculation,” on the hypothesis that psychological resistance against deceitful persuasion attempts can be conferred by exposing people to a weakened version of the misinformation. Inoculation theory has been extensively studied in the context of building resistance to misinformation in particular domains, such as climate science (Banas and Rains, 2010; Compton, 2013) and vaccines (Jolley and Douglas, 2017), but only recently has research begun to examine the effectiveness of “broad base” misinformation inoculation. A series of studies investigated this in the context of a recently developed fake news game (Basol et al., 2020; Roozenbeek and Linden, 2019). In this game players take the role of a news editor whose goal is to build a media empire using fake news articles. In playing this game, participants are exposed to “weakened doses” of misinformation techniques with the aim that this will improve their ability to spot misinformation techniques and, thus, they will become “inoculated” against misinformation. These studies found a significant reduction in

the perceived accuracy of fake news presented in the form of Twitter “tweets” irrespective of education, age, political ideology, and cognitive style.

A further avenue worth exploration is utilization of the wisdom of crowds. A recent study found that crowdsourced judgments of news source quality by laypeople correlated strongly with the judgments of professional fact checkers (Pennycook and Rand, 2019c). If this result is reliable, the wisdom of crowds might be harnessed by algorithms on social media to identify reliable news and make it much more visible than less reliable news.

Summary.

In the interests of space, the survey of work on the cognitive science of fake news has been brief and partial. It leaves us with many open questions. To that extent, however, it is an accurate reflection of the state of our knowledge here: many central questions remain for future research. To mention just two: (1) to what extent does analytic thinking (which appears to predict the capacity to distinguish fake news from real) reduce the propensity to *share* fake news? Intuitively, reduced acceptance would predict a reduced likelihood of sharing, but this remains controversial, with Pennycook and colleagues finding the expected relation (e.g. (Pennycook et al., 2020b) but others failing to find an association (e.g. Osmundsen et al., 2020). (2) to what extent do these lab-based experiments capture real world behaviour? There is much greater need for researchers in different disciplines and subdisciplines to take account of one another’s work: too few take into account the prevalence of expressive responding and trolling, for example. We often decry the extent to which people are siloed into filter bubbles on the internet, but researchers may need to burst bubbles of their own.⁴

Notes

¹ Note that cognitive scientists tend to define fake news more narrowly (the single most cited paper on the topic Lazer et al. (2018) defines it as “fabricated information that mimics news media content in form but not in organizational process or intent” (1094). We follow philosophers like Mukerji (2018) in encompassing ‘bullshit’ information – produced by those indifferent to truth – as well as information produced with the intention to deceive.

² In addition to the use of motivated inference and expressing responding, survey respondents may simply fail to respond sincerely for fun or to troll researchers (Lopez and Hillygus, 2018). Our own unpublished research indicates that this kind of response is surprisingly common.

³ We acknowledge that there are alternative explanations available. Public affirmations of belief, especially with high costs attached, provide an opportunity for an especially powerful expression of one’s attitudes, along similar lines to the ways in which costs can enhance the strength of signals (Sosis, 2003) or increase credibility (Henrich, 2009). Arguably, the states that underlie such a signal could count as beliefs (Funkhouser, 2017), but this remains an open question. Since the states play many of the same functional roles as beliefs – in particular, generating behaviour of much the same kind as the belief would – we take them to raise many of the same questions and therefore suggest that we need not wait until their doxastic status has been settled to take them as an important target.

⁴ We are grateful to the editors of this volume and a reviewer for helpful comments. We acknowledge the generous support of the Australian Research Council (DP180102384).

Notes on Contributors

Neil Levy is professor of philosophy at Macquarie University and a senior research fellow at the Uehiro Centre for Practical Ethics, University of Oxford. He works on many topics, with an emphasis on issues at the intersection of philosophy and psychology.

Robert M Ross is a post-doctoral research fellow in the Department of Philosophy at Macquarie University. His work focuses on the cognitive science of belief.

References

- Allen, J., Howland, B., Mobius, M., Rothschild, D., Watts, D.J., 2020. Evaluating the fake news problem at the scale of the information ecosystem. *Science Advances* 6, eaay3539. <https://doi.org/10.1126/sciadv.aay3539>
- Bago, B., Rand, D.G., Pennycook, G., 2020. Fake news, fast and slow: Deliberation reduces belief in false (but not true) news headlines. *J Exp Psychol Gen.* <https://doi.org/10.1037/xge0000729>
- Basol, M., Roozenbeek, J., van der Linden, S., 2020. Good News about Bad News: Gamified Inoculation Boosts Confidence and Cognitive Immunity Against Fake News. *J Cogn* 3, 2. <https://doi.org/10.5334/joc.91>
- Berinsky, A.J., 2018. Telling the Truth about Believing the Lies? Evidence for the Limited Prevalence of Expressive Survey Responding. *The Journal of Politics* 80, 211–224. <https://doi.org/10.1086/694258>
- Bullock, J.G., Gerber, A.S., Hill, S.J., Huber, G.A., 2015. Partisan Bias in Factual Beliefs about Politics. *QJPS* 10, 519–578. <https://doi.org/10.1561/100.00014074>
- Bullock, J.G., Lenz, G., 2019. Partisan Bias in Surveys. *Annual Review of Political Science* 22, 325–342. <https://doi.org/10.1146/annurev-polisci-051117-050904>
- Chan, M.-P.S., Jones, C.R., Hall Jamieson, K., Albarracín, D., 2017. Debunking: A Meta-Analysis of the Psychological Efficacy of Messages Countering Misinformation. *Psychol Sci* 28, 1531–1546. <https://doi.org/10.1177/0956797617714579>
- Clayton, K., Blair, S., Busam, J.A., Forstner, S., Glance, J., Green, G., Kawata, A., Kovvuri, A., Martin, J., Morgan, E., Sandhu, M., Sang, R., Scholz-Bright, R., Welch, A.T., Wolff, A.G., Zhou, A., Nyhan, B., 2019. Real Solutions for Fake News? Measuring the Effectiveness of General Warnings and Fact-Check Tags in Reducing Belief in False Stories on Social Media. *Polit Behav.* <https://doi.org/10.1007/s11109-019-09533-0>
- Clifford, S., Kim, Y., Sullivan, B.W., 2020. An Improved Question Format for Measuring Conspiracy Beliefs. *Public Opin Q.* <https://doi.org/10.1093/poq/nfz049>
- Dias, N., Pennycook, G., Rand, D.G., 2020. Emphasizing publishers does not effectively reduce susceptibility to misinformation on social media. *Harvard Kennedy School Misinformation Review* 1. <https://doi.org/10.37016/mr-2020-001>
- Dunning, D., 2019. Gullible to ourselves, in: *The Social Psychology of Gullibility: Conspiracy Theories, Fake News and Irrational Beliefs*. Routledge, Abingdon, pp. 217–233.
- Freeze, M., Baumgartner, M., Bruno, P., Gunderson, J.R., Olin, J., Ross, M.Q., Szafran, J., 2020. Fake Claims of Fake News: Political Misinformation, Warnings, and the Tainted Truth Effect. *Polit Behav.* <https://doi.org/10.1007/s11109-020-09597-3>
- Funkhouser, E., 2017. Beliefs as Signals: A New Function for Belief. *Philosophical Psychology* 30, 809–831. <https://doi.org/10.1080/09515089.2017.1291929>
- Grice, 2017. Fake news handed Brexiteers the referendum – and now they have no idea what they’re doing. *The Independent*.
- Guess, A.M., Lerner, M., Lyons, B., Montgomery, J.M., Nyhan, B., Reifler, J., Sircar, N., 2020a. A digital media literacy intervention increases discernment between mainstream and false news in the United States and India. *PNAS* 117, 15536–15545. <https://doi.org/10.1073/pnas.1920498117>
- Guess, A.M., Lockett, D., Lyons, B., Montgomery, J.M., Nyhan, B., Reifler, J., 2020b. “Fake news” may have limited effects beyond increasing beliefs in false claims. *Harvard Kennedy School Misinformation Review* 1. <https://doi.org/10.37016/mr-2020-004>

- Guess, A.M., Nyhan, B., Reifler, J., 2020c. Exposure to untrustworthy websites in the 2016 US election. *Nature Human Behaviour* 4, 472–480. <https://doi.org/10.1038/s41562-020-0833-x>
- Gunther, R., Beck, P.A., Nisbet, E.C., 2018. Fake news did have a significant impact on the vote in the 2016 election: Original full-length version with methodological appendix.
- Hannon, M., 2019. Political Disagreement or Partisan Cheerleading? The Role of Expressive Discourse in Politics.
- Henrich, J., 2009. The evolution of costly displays, cooperation and religion: credibility enhancing displays and their implications for cultural evolution. *Evolution and Human Behavior* 30, 244–260. <https://doi.org/10.1016/j.evolhumbehav.2009.03.005>
- Jerit, J., Barabas, J., 2012. Partisan Perceptual Bias and the Information Environment. *The Journal of Politics* 74, 672–684. <https://doi.org/10.1017/S0022381612000187>
- Jolley, D., Douglas, K.M., 2017. Prevention is better than cure: Addressing anti-vaccine conspiracy theories. *Journal of Applied Social Psychology* 47, 459–469.
- Kahan, D.M., 2017. Misconceptions, Misinformation, and the Logic of Identity-Protective Cognition (SSRN Scholarly Paper No. ID 2973067). Social Science Research Network, Rochester, NY.
- Kahan, D.M., 2016. The Politically Motivated Reasoning Paradigm, Part 1: What Politically Motivated Reasoning Is and How to Measure It - Kahan - - Major Reference Works - Wiley Online Library, in: *Emerging Trends in Social and Behavioral Sciences*. pp. 1-16.
- Kahan, D.M., 2015. Climate-Science Communication and the Measurement Problem. *Political Psychology* 36, 1–43.
- Lazer, D.M.J., Baum, M.A., Benkler, Y., Berinsky, A.J., Greenhill, K.M., Menczer, F., Metzger, M.J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S.A., Sunstein, C.R., Thorson, E.A., Watts, D.J., Zittrain, J.L., 2018. The science of fake news. *Science* 359, 1094–1096. <https://doi.org/10.1126/science.aao2998>
- Lerman, A.E., Sadin, M.L., Trachtman, S., 2017. Policy Uptake as Political Behavior: Evidence from the Affordable Care Act. *American Political Science Review* 111, 755–770. <https://doi.org/10.1017/S0003055417000272>
- Levy, N., 2017. The Bad News About Fake News 6, 18.
- Lewandowsky, S., Ecker, U.K.H., Seifer, C.M., Schwarz, N., Cook, J., 2012. Misinformation and Its Correction: Continued Influence and Successful Debiasing. *Psychological Science in the Public Interest* 13, 106–131.
- Lopez, J., Hillygus, D.S., 2018. Why So Serious?: Survey Trolls and Misinformation (SSRN Scholarly Paper No. ID 3131087). Social Science Research Network, Rochester, NY. <https://doi.org/10.2139/ssrn.3131087>
- Lutzke, L., Drummond, C., Slovic, P., Árvai, J., 2019. Priming critical thinking: Simple interventions limit the influence of fake news about climate change on Facebook. *Global Environmental Change* 58, 101964. <https://doi.org/10.1016/j.gloenvcha.2019.101964>
- McGrath, M.C., 2017. Economic Behavior and the Partisan Perceptual Screen. *QJPS* 11, 363–383. <https://doi.org/10.1561/100.00015100>
- Mercier, H., 2020. *Not Born Yesterday: The Science of Who We Trust and What We Believe*. Princeton University Press, Princeton.
- Merkley, E., 2020. Anti-Intellectualism, Populism, And Motivated Resistance To Expert Consensus. *Public Opin Q.* <https://doi.org/10.1093/poq/nfz053>

- Mukerji, N., 2018. What is Fake News? *Ergo: An Open Access Journal of Philosophy* 5, 923–946. <https://doi.org/10.3998/ergo.12405314.0005.035>
- Nyhan, B., Reifler, J., 2010. When Corrections Fail: The Persistence of Political Misperceptions. *Polit Behav* 32, 303–330. <https://doi.org/10.1007/s11109-010-9112-2>
- Osmundsen, M., Bor, A., Vahlstrup, P.B., Bechmann, A., Petersen, M.B., 2020. Partisan polarization is the primary psychological motivation behind “fake news” sharing on Twitter (preprint). *PsyArXiv*. <https://doi.org/10.31234/osf.io/v45bk>
- Pennycook, G., Bear, A., Collins, E., 2020a. The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Headlines Increases Perceived Accuracy of Headlines Without Warnings. *Management Science*. <https://doi.org/10.1287/mnsc.2019.3478>
- Pennycook, G., Epstein, Z., Mosleh, M., Arechar, A.A., Eckles, D., Rand, D., 2019. Understanding and reducing the spread of misinformation online. <https://doi.org/10.31234/osf.io/3n9u8>
- Pennycook, G., McPhetres, J., Zhang, Y., Rand, D., 2020b. Fighting COVID-19 misinformation on social media: Experimental evidence for a scalable accuracy nudge intervention. *Psychological Science* 31, 770–780. <https://doi.org/10.31234/osf.io/uwbk9>
- Pennycook, G., Rand, D.G., 2019a. Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition, The Cognitive Science of Political Thought* 188, 39–50. <https://doi.org/10.1016/j.cognition.2018.06.011>
- Pennycook, G., Rand, D.G., 2019b. Cognitive Reflection and the 2016 U.S. Presidential Election. *Pers Soc Psychol Bull* 45, 224–239. <https://doi.org/10.1177/0146167218783192>
- Pennycook, G., Rand, D.G., 2019c. Fighting misinformation on social media using crowdsourced judgments of news source quality. *PNAS* 116, 2521–2526. <https://doi.org/10.1073/pnas.1806781116>
- Peterson, E., Iyengar, S., 2020. Partisan Gaps in Political Information and Information-Seeking Behavior: Motivated Reasoning or Cheerleading? *American Journal of Political Science* n/a. <https://doi.org/10.1111/ajps.12535>
- Porter, E., Wood, T., Kirby, D., 2018. Sex Trafficking, Russian Infiltration, Birth Certificates, and Pedophilia: A Survey Experiment Correcting Fake News. *Journal of Experimental Political Science* 5, 1–6. <https://doi.org/10.1017/XPS.2017.32>
- Porter, E., Wood, T.J., 2019. *False Alarm: The Truth about Political Mistruths in the Trump Era*. Cambridge University Press, S.I.
- Prior, M., Sood, G., Khanna, K., 2015. You Cannot be Serious: The Impact of Accuracy Incentives on Partisan Bias in Reports of Economic Perceptions. *QJPS* 10, 489–518. <https://doi.org/10.1561/100.00014127>
- Raphelson, S., 2018. Survivors Of Mass Shootings Face Renewed Trauma From Conspiracy Theorists [WWW Document]. *NPR.org*. URL <https://www.npr.org/2018/03/20/595213740/survivors-of-mass-shootings-face-renewed-trauma-from-conspiracy-theorists> (accessed 7.3.19).
- Robles, F., 2016. Florida Woman Is Charged With Threatening Sandy Hook Parent. *The New York Times*.
- Roozenbeek, J., Linden, S. van der, 2019. Fake news game confers psychological resistance against online misinformation. *Palgrave Commun* 5, 1–10. <https://doi.org/10.1057/s41599-019-0279-9>

- Rose, D., Buckwalter, W., Turri, J., 2014. When Words Speak Louder Than Actions: Delusion, Belief, and the Power of Assertion. *Australasian Journal of Philosophy* 1–18. <https://doi.org/10.1080/00048402.2014.909859>
- Roush, C.E., Sood, G., 2020. A Gap in Our Understanding? Reconsidering the Evidence for Partisan Knowledge Gaps. Presented at the Annual Meeting of the Southern Political Science Association., p. 73.
- Rust, J., Schwitzgebel, E., 2009. Ethicists' and Nonethicists' Responsiveness to Student E-mails: Relationships Among Expressed Normative Attitude, Self-Described Behavior, and Empirically Observed Behavior. *Philosophical Psychology* 22, 711–725.
- Schaffner, B.F., Luks, S., 2018. Misinformation Or Expressive Responding? What An Inauguration Crowd Can Tell Us About The Source Of Political Misinformation In Surveys. *Political Opinion Quarterly* 82, 135–147.
- Schwitzgebel, E., 2009. Do Ethicists Steal More Books? *Philosophical Psychology* 22, 711–725. <https://doi.org/10.1080/09515080903409952>
- Sosis, R., 2003. Why aren't we all hutterites? *Hum Nat* 14, 91–127. <https://doi.org/10.1007/s12110-003-1000-6>
- Swire-Thompson, B., DeGutis, J., Lazer, D., 2020. Searching for the backfire effect: Measurement and design considerations (preprint). *PsyArXiv*. <https://doi.org/10.31234/osf.io/ba2kc>
- Taber, C.S., Lodge, M., 2006. Motivated Skepticism in the Evaluation of Political Beliefs. *American Journal of Political Science* 50, 755–769. <https://doi.org/10.1111/j.1540-5907.2006.00214.x>
- Tappin, B.M., Pennycook, G., Rand, D.G., 2020a. Bayesian or biased? Analytic thinking and political belief updating. *Cognition* 204, 104375. <https://doi.org/10.1016/j.cognition.2020.104375>
- Tappin, B.M., Pennycook, G., Rand, D.G., 2020b. Thinking clearly about causal inferences of politically motivated reasoning: why paradigmatic study designs often undermine causal inference. *Current Opinion in Behavioral Sciences* 34, 81–87. <https://doi.org/10.1016/j.cobeha.2020.01.003>
- Tetlock, P.E., 2003. Thinking the unthinkable: sacred values and taboo cognitions. *Trends in Cognitive Sciences* 7, 320–324. [https://doi.org/10.1016/S1364-6613\(03\)00135-9](https://doi.org/10.1016/S1364-6613(03)00135-9)
- Uscinski, J.E., Parent, J.M., 2014. *American Conspiracy Theories*. Oxford University Press.
- Vosoughi, S., Roy, D., Aral, S., 2018. The spread of true and false news online. *Science* 359, 1146–1151. <https://doi.org/10.1126/science.aap9559>
- Walter, N., Cohen, J., Holbert, R.L., Morag, Y., 2019. Fact-Checking: A Meta-Analysis of What Works and for Whom. *Political Communication*. <https://doi.org/10.1080/10584609.2019.1668894>
- Walter, N., Tukachinsky, R., 2020. A Meta-Analytic Examination of the Continued Influence of Misinformation in the Face of Correction: How Powerful Is It, Why Does It Happen, and How to Stop It? <https://doi.org/10.1177/0093650219854600>
- Wood, T., Porter, E., 2019. The Elusive Backfire Effect: Mass Attitudes' Steadfast Factual Adherence. *Polit Behav* 41, 135–163. <https://doi.org/10.1007/s11109-018-9443-y>
- Zaller, J., Feldman, S., 1992. A Simple Theory of the Survey Response: Answering Questions versus Revealing Preferences. *American Journal of Political Science* 36, 579. <https://doi.org/10.2307/2111583>