

Review

The Psychology of Fake News

Gordon Pennycook^{1,2,*} and David G. Rand^{3,4,5,*}

We synthesize a burgeoning literature investigating why people believe and share false or highly misleading news online. Contrary to a common narrative whereby politics drives susceptibility to fake news, people are ‘better’ at discerning truth from falsehood (despite greater overall belief) when evaluating politically concordant news. Instead, poor truth discernment is associated with lack of careful reasoning and relevant knowledge, and the use of heuristics such as familiarity. Furthermore, there is a substantial disconnect between what people believe and what they share on social media. This dissociation is largely driven by inattention, more so than by purposeful sharing of misinformation. Thus, interventions can successfully nudge social media users to focus more on accuracy. Crowdsourced veracity ratings can also be leveraged to improve social media ranking algorithms.

Toward a Psychology of False and Misleading Online News

Fabricated news is nothing new. For example, in 1835 *The Sun* newspaper in New York published six articles about purported life on the moon which came to be known as the ‘Great Moon Hoax’. During the 2016 US Presidential Election and UK Brexit Referendum, however, a different form of **fake news** (see [Glossary](#)) rose to prominence ([Box 1](#)): false or highly misleading political ‘news’ stories, primarily originating on social media [1]. Concern about fake news was redoubled in 2020 in the face of widespread **misinformation** and **disinformation** [2] on social media about the coronavirus disease 2019 (COVID-19) pandemic [3] and the 2020 US Presidential Election [4]. Misleading **hyperpartisan news**, as well as **yellow journalism** [5], are related forms of problematic news content that are likely sources of political polarization [6]. What is it about human psychology – and its interaction with social media [7,8] – that explains the failure to distinguish between accurate and inaccurate content online? Apart from being of theoretical interest, this question has practical consequences: developing effective interventions against misinformation depends on understanding the underlying psychology.

We focus here primarily on online content that is presented in the form of news articles. However, false and misleading claims come in many forms, and there are several literatures that are clearly related, but outside the scope of our review (although we will draw some connections throughout). These include work on conspiracy belief [9], superstition [10], rumors [11], bullshit receptivity [12], and misperceptions [13], among others. Furthermore, our focus is on individual examples of misinformation, and not on organized disinformation campaigns (e.g., by the Russian Internet Research Agency, or campaigns relating to global warming or fraud in the 2020 US Presidential Election).

Why Do People Fall for Fake News?

When considering the factors that may influence what people believe, it is essential to distinguish between two fundamentally different ways to conceptualize belief in true and false news. One common approach is to focus on truth ‘discernment’, or the extent to which misinformation is believed ‘relative’ to accurate content. Discernment, typically calculated as belief in true news minus belief in false news (akin to ‘sensitivity’ or *d'* in signal detection theory [14]) captures the ‘overall’ accuracy of one’s beliefs – and thus gives insight into failures to distinguish between true and false content (‘falling for fake news’).

Highlights

Recent evidence contradicts the common narrative that partisanship and politically motivated reasoning explain why people fall for ‘fake news’.

Poor truth discernment is linked to a lack of careful reasoning and relevant knowledge, as well as to the use of familiarity and source heuristics.

There is also a large disconnect between what people believe and what they will share on social media, and this is largely driven by inattention rather than by purposeful sharing of misinformation.

Effective interventions can nudge social media users to think about accuracy, and can leverage crowdsourced veracity ratings to improve social media ranking algorithms.

¹Hill/Levene Schools of Business, University of Regina, Regina, SK S4S 0A2, Canada

²Department of Psychology, University of Regina, Regina, SK S4S 0A2, Canada

³Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02142, USA

⁴Institute for Data, Systems, and Society, Massachusetts Institute of Technology, Cambridge, MA 02142, USA

⁵Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02142, USA

*Correspondence: gordon.pennycook@uregina.ca (G. Pennycook) and drand@mit.edu (D.G. Rand).



Box 1. Prevalence of Fake News

Various analyses of social media and web browsing data have been used in an attempt to determine the prevalence of fake news, often with a focus on the 2016 US Presidential Election. For example, using web browsing data, archives of fact-checking websites, and a survey, Allcott and Gentzkow [19] estimated that a particular set of news stories that are known to be false were shared on Facebook at least 38 million times in the 3 months leading up to the 2016 election (30 million of which were for news favoring Donald Trump). This estimate represents a lower bound since it only reflects that specific set of known false news.

Other analyses have focused on fake news publishers (i.e., websites) rather than on individual articles. Based on data from Twitter [117], Facebook [77, 118], and web browsing [89], these studies concluded that content from known fake news sites represents a small proportion of most people's media diets, and that the average social media user was exposed to little fake news during the 2016 election.

These analyses have important limitations, however, because the only available data concern what people are sharing and what they visit when they click through to visit news sites off-platform. But, of course, the vast majority of the time that people are exposed to news on social media, they simply read the post without sharing it or clicking on the link to visit the actual source website. Furthermore, so-called 'fake news' only represents one category of misinformation, and misleading content from sources such as hyperpartisan news websites likely represents a much larger proportion of people's media diets [6, 119]. Thus, the actual on-platform exposure of the average user to misinformation remains an open question [120]. We feel it is premature to conclude that exposure rates are minimal, and thus that false and misleading news online is not a problem (also [7, 8]). This is especially true when looking beyond the 2016 election because new misinformation threats – such as false claims about COVID-19 [3, 44] and fraud in the 2020 US Presidential Election [4] – have gained widespread traction through amplification by (mostly Republican) political elites.

Accordingly, exposure to fake news (and misinformation more broadly) is not equally distributed across all users. In particular, political conservatives and older adults were far more likely to visit fake news websites or share fake news articles during the 2016 Presidential Election [19, 89, 117, 118]. Studies have also found associations between political conservatism and belief in misinformation in the USA [20, 44], Chile [121], and Germany [122], but not in Hungary [24], and users who engage in less reasoning have been found to share content from lower-quality news sites on Twitter [71]. Thus, even if it was true that the average social media user was not exposed to that much misinformation, exposures rates are substantially higher in subpopulations that may be particularly vulnerable to believing inaccurate content. Finally, misinformation that originates on social media sometimes transitions to much larger audiences when it is picked up by traditional media outlets – either via direct repetition or debunking (which may result in inadvertent amplification).

Another approach is to focus on overall belief, or the extent to which news – regardless of its accuracy – is believed (calculated as the average or sum of belief in true news and belief in false news, akin to calculating 'bias' in signal detection theory [14]). Critically, factors that alter overall belief need not impact people's ability to tell truth from falsehood [15]: increasing or decreasing belief in true and false headlines to an equivalent extent has no effect on the overall accuracy of one's beliefs (i.e., does not affect truth discernment).

Political Motivations

A popular narrative is that the failure to discern between true and false news is rooted in political motivations. For example, it has been argued that people are motivated consumers of (mis)information [16] – that they engage in 'identity-protective cognition' when faced with politically valenced content, and this leads them to be overly believing of content that is consistent with their partisan identity and overly skeptical of content that is inconsistent with their partisan identity [17]. A related theory argues that people place loyalty to their political identities above the truth – and thus fail to discern truth from falsehood in favor of simply believing ideologically concordant information [18]. These accounts contend that a strong causal influence of political motivation on belief is thus the dominant factor explaining why people fall for fake news.

It is clearly true that partisanship is associated with overall belief: People are more likely to believe news content that is concordant (versus discordant) with their political partisanship [19–25] (Figure 1B). It is important to note, however, that the effect of political concordance is typically much smaller than that of the actual veracity of the news [20, 21, 26]. In other words, true but

Glossary

Algorithm: steps or calculations performed by a computer (given a set of rules) to solve a problem or complete a task. In the context of social media, algorithms are used to determine what content users see.

Disinformation: information that is false or inaccurate, and that was created with a deliberate intention to mislead people.

Fake news: news content published on the internet that aesthetically resembles actual legitimate mainstream news content, but that is fabricated or extremely inaccurate. Also referred to as false, junk, or fabricated news.

Hyperpartisan news: news content that is not entirely fabricated, but which covers events that actually occurred with a strong partisan bias. As a result, hyperpartisan news is typically misleading, and we therefore include it as a form of misinformation.

Misinformation: information that is false, inaccurate, or misleading. Unlike disinformation, misinformation does not necessarily need to be created deliberately to mislead. Misinformation is sometimes used to refer exclusively to inaccuracies that are accidental; however, because it is difficult to ascertain the intentions of the unknown individuals who create falsehoods that spread on the internet, we use misinformation as a broader umbrella term here (i.e., much of the content used in the studies we discuss could be classified as disinformation and/or hyperpartisan news and/or propaganda, etc.).

Yellow journalism: content from newspapers, magazines, or websites that is poorly researched and sensationalist, and that is created with the goal of increasing sales or, on the internet, clicks. Roughly equivalent to tabloid journalism.

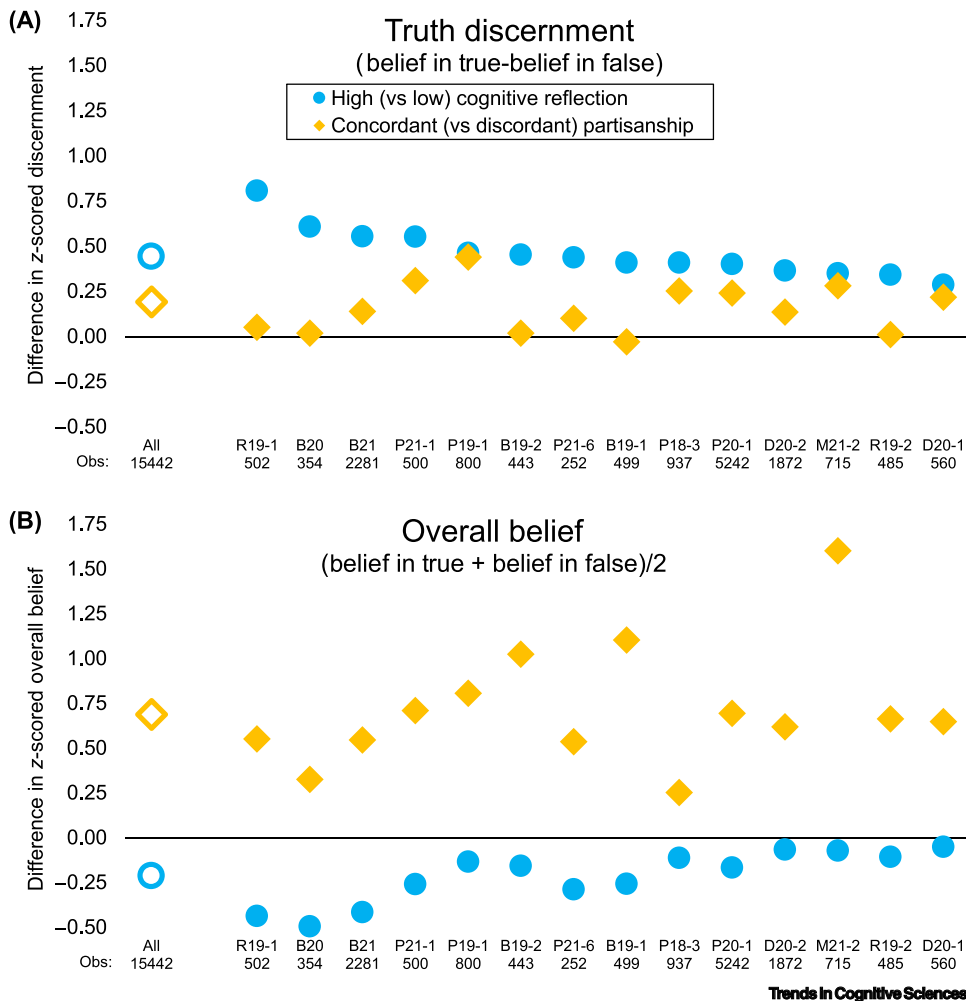


Figure 1. Political Concordance and Cognitive Reflection Have Divergent Relationships with Belief. Reasoning (or lack thereof) is much more strongly related to truth discernment than is political concordance, whereas political concordance is much more strongly related to overall belief than is reasoning. Data from 14 studies [20,21,26,34,35,41,49,61,84,93] in which US participants were asked to rate the accuracy of actual news posts from social media (headlines equally balanced true versus false, as well as Democrat-consistent versus Republican-consistent). Study names are indicated by the first author's last initial and publication date, followed by a dash and study number for multistudy papers; the number of participants in each study is indicated under the name. Data from all 14 studies are pooled to calculate the 'All' values. (A) Difference in truth discernment (belief in true news minus belief in false news) between participants who scored higher on the Cognitive Reflection Test relative to those who scored lower (median split; collapsing over headline political concordance) and between headlines that were politically concordant relative to politically discordant (collapsing over participant cognitive reflection). More-reflective participants were more discerning than less-reflective participants, and discernment was higher for politically concordant headlines than for politically discordant headlines. Critically, the difference based on reasoning was roughly twice as large as the difference based on partisanship. (B) Difference in overall belief (average of belief in true news and belief in false news) between participants who scored higher on the Cognitive Reflection Test relative to those who scored lower (median split; collapsing over headline political concordance), and between headlines that were politically concordant relative to politically discordant (collapsing over participant cognitive reflection). More-reflective participants had a somewhat lower overall belief than less-reflective participants, whereas overall belief was higher for politically concordant headlines than for politically discordant headlines. Abbreviation: Obs., observed.

politically discordant news is typically believed much more than false but politically concordant news – politics does not trump truth. Furthermore, greater overall belief in politically consistent news does not necessarily indicate politically motivated reasoning. Such differences could even arise from unbiased rational (e.g., Bayesian) inference built on prior factual beliefs that

differ across party lines (e.g., owing to exposure to different information environments) [27–33] (Box 2 for details).

We now turn to the impact of political concordance on truth discernment. Greater overall belief in politically concordant news might seem to suggest that people are more inaccurate when assessing politically concordant news – in other words, that political concordance (and the associated motivations) interferes with truth discernment. In fact, the data reveal the opposite pattern: People are somewhat better at discerning truth from falsehood when judging politically concordant news compared with politically discordant news (Figure 1A). Taken together, the evidence therefore suggests that political identity and politically motivated reasoning are not the primary factors driving the inability to tell truth from falsehood in online news.

Reasoning

Another perspective on the (in)ability to differentiate between truth and falsehood comes from the field of reasoning. Work in this vein has a particular focus on dual-process theories stipulating that analytic thinking can override automatic, intuitive responses (Box 3 for details). The key question this perspective asks is – what is the role of reflective reasoning in the ability to discern fake news from truth?

One potential answer, which follows from the earlier-referenced work on political identity, argues that deliberative ('System 2') reasoning is often motivated by political identity, and that people

Box 2. Challenges in Identifying Politically Motivated Reasoning

The observation that people are more likely to believe information that is consistent with their political ideology/partisanship (and are less likely to believe information that is inconsistent with their ideology/partisanship) is often taken as evidence for politically motivated reasoning [22, 123, 124]. Critically, however, this pattern does not actually provide clear evidence of politically motivated reasoning because partisan identity is likely confounded with other relevant variables [27, 125]. Most notably, partisans differ in what they believe about the world, even when it comes to so-called 'factual beliefs' – beliefs that relate to facts or empirical evidence [28], such as in the case of global warming [126]. This is critical because a large body of evidence from entirely nonpolitical contexts shows that what people believe to be true about the world influences their reasoning (a phenomenon known as belief bias [127]). Indeed, once prior beliefs are accounted for, the apparent impact of political concordance on processes of belief formation is typically greatly reduced or eliminated [28, 29, 128]. Thus, observing a difference across ideological lines is not sufficient evidence to conclude that partisan identity or political motivations are themselves responsible for the difference [27, 125].

To clearly disentangle the impact of partisan identity or motivations versus prior beliefs, further experimental work is needed – for example, studies that manipulate prior factual beliefs and/or political motivations [27, 30]. Relatedly, understanding the origins of partisan differences in prior factual beliefs is also of the utmost importance. Exposure to different information streams is a promising candidate: entirely rational (e.g., Bayesian) and truth-seeking (i.e., nonpolitically motivated) people who obtain their information from conservative (e.g., Fox News) versus liberal (e.g., MSNBC) news sources would naturally wind up with very different factual beliefs about the world. Evaluating new information in light of how well it aligns with one's priors – although often referred to as 'confirmation bias' – is not, in fact, necessarily evidence of bias in a normative sense. When there is uncertainty about the reliability of information sources or data-generating processes, it can be entirely consistent with Bayesian inference to be skeptical of information that is inconsistent with one's prior factual beliefs [30, 129]. In such cases, Bayesian agents may infer that the source or the data-generating process is unreliable rather than that their prior belief was incorrect – and doing so would not be evidence of bias *per se* (where bias is defined as deviating from some normative, e.g., Bayesian, benchmark) [130].

Finally, the critique described in this box also applies to inferring 'motivated System 2 reasoning' from the observation that cognitive sophistication is sometimes (but not always or even frequently [131–133]) associated with polarization rather than accuracy [17]. The association between cognitive sophistication and polarization (e.g., in the context of climate change) disappears entirely once prior factual beliefs are accounted for [28]. Instead, it appears that more cognitively sophisticated individuals may place more weight on their prior factual beliefs when evaluating new evidence, instead of placing more weight on concordance with their political identities. Of course, people's prior factual beliefs may themselves be caused by politically motivated reasoning – but they need not be, and most study designs cannot determine (but must assume) a causal connection.

Box 3. Dual-Process Models of Reasoning and the Consequences of Deliberation

Dual-process theories are a core component of research on the cognitive science of reasoning. These theories argue that human cognition can be partitioned into two fundamentally different types of processes that differ in terms of their characteristics [40,134,135]: Type 1 (or System 1) processing that is characterized primarily by automaticity such that Type 1 outputs ('intuitions') come to mind directly as a response to the stimulus, and Type 2 (or System 2) processing that is characterized by the deliberation that may or may not arise given a particular intuitive output (or set of outputs).

Consider the following problem from the Cognitive Reflection Test [136]: 'A bat and ball cost \$1.10 in total. The bat costs \$1.00 more than the ball. How much does the ball cost?' The incorrect intuitive answer – 10 cents – comes to mind intuitively for most people, whereas the correct answer – 5 cents – emerges (for most) only with an additional deliberation process [137]. Performance on tasks such as the Cognitive Reflection Test are associated with a wide range of beliefs and behaviors [138], including actual social media behaviors observed on Twitter ([71]; also [109]), as well as conspiracy ideation [139], politics [140], bullshit receptivity [12], and endorsement of a variety of epistemically suspect beliefs [138]. Importantly, although dual-process theories typically emphasize the importance of overriding incorrect intuitions via analytic thinking, this should not be taken to imply that intuitions are always incorrect or that analytic thinking is always accurate [135].

engage in 'identity protective cognition' [16]. This account predicts that engaging in more deliberation should lead to more politically polarized beliefs – and, most importantly, to greater belief in politically concordant but false claims. Thus, more deliberation should be associated with worse truth discernment.

By contrast, more 'classical' reasoning accounts (i.e., accounts that are more consistent with work on dual-process reasoning in other domains) portray System 2 reasoning as being responsible for correcting faulty intuitions (Box 3). This perspective therefore predicts that people who deliberate more will simply be less likely to believe false content – and are better able to discern between true and false content – regardless of the political concordance of the news that they are evaluating.

Across numerous recent studies, the evidence supports the classical reasoning account over the motivated System 2 reasoning account. People who are more reflective (Box 3) are less likely to believe false news content – and are better at discerning between truth and falsehood – regardless of whether the news is consistent or inconsistent with their partisanship [20,26,34,35] (Figure 1A). The same pattern is evident with respect to discernment between biased and misleading hyperpartisan news and true (mainstream) news [35] and when judging full news stories as opposed to just headlines [36], and using measures beyond the Cognitive Reflection Test (Figure 1), such as thinking disposition questionnaires [34] and the Berlin Numeracy Test [35]. Belief in fake news is also associated with delusionality [34], dogmatism [34], religious fundamentalism [34], bullshit receptivity [37], and overclaiming [37] (all factors associated with analytic thinking; Box 3). Furthermore, experimentally manipulating participants' level of deliberation demonstrates a causal effect whereby deliberation reduces belief in false (but not true) news, regardless of partisan alignment (and has no effect on polarization) [26]. Research also shows that overconfidence may contribute to susceptibility to false information [38], perhaps because it stops people from slowing down and engaging in reflective reasoning [39,40].

How, then, do people determine news veracity? The correlation between cognitive reflection and disbelief in fake news is stronger in cases where the content is more obviously implausible (and vice versa for true news) [20]. This suggests that, in cases where people actually do stop and think, relevant prior knowledge is likely to be a critical factor. Indeed, political knowledge is positively associated truth discernment for political news content [23,41], as is media literacy [42] and general information literacy [43]. Similarly, basic science knowledge is positively associated with truth discernment for (mis)information about COVID-19 [44]. This implies, unfortunately, that reasoning may not improve accuracy in contexts where prior knowledge is heavily

distorted (e.g., by partisan media consumption or misinformation campaigns by political elites – climate change being a prime example) [28] (Box 2).

Thus, when it comes to the role of reasoning, it seems that people fail to discern truth from falsehood because they do not stop to reflect sufficiently on their prior knowledge (or have insufficient or inaccurate prior knowledge) – and not because their reasoning abilities are hijacked by political motivations.

Heuristics

Prior work in judgment and decision making [45] indicates that people are likely to use heuristics or mental shortcuts when judging news headlines. What, then, are the specific features of fake news that influence people's intuitions or cause them to make mistakes when reasoning?

One key route to intuitive belief in news is familiarity [37]. The influence of prior exposure on judgments of truth – sometimes referred to as the illusory truth effect – is well documented [46,47]. Indeed, wartime rumors during WWII that were more familiar were more likely to be believed [48]. Consistent with this, a single prior exposure to a fake news headline increases later belief in the headline [49,50]. Remarkably, this is evident even if the headline is extremely implausible (also [51]) and inconsistent with one's political partisanship [49]. Thus, feelings of familiarity and, possibly, processing fluency *per se* ([52,53]; (but see [54]) likely contributes to increased belief in false claims.

The source is another important cue that may be used when evaluating news. Participants are more likely to believe information provided by people whom they view as being credible (reviewed in [55]), and a large literature from political science has robustly demonstrated the impact of elite messaging, in particular, on public opinion [56]. For example, attributing a false claim to President Trump increased Trump supporters' belief in the claim while reducing Democrats' belief in the claim [57]. Furthermore, social feedback provided by social media platforms (e.g., 'likes') also increases belief in news content, particularly for misinformation [58] – a factor that is likely to covary with elite messaging (given that political elites, such as Presidents and people in Congress or the media, often have many social media followers).

Finally, a salient feature of fake news headlines also seems to be that they are often emotionally evocative. That is, fake news is often geared toward provoking shock, fear, anger [8,59], or (more broadly) moral outrage [60]. This is important because people who report experiencing more emotion (positive or negative) at the outset of the task are more likely to believe false (but not true) news; and instructing people to rely on emotion increases belief in false (but not true) headlines [61].

Believing versus Sharing Fake News

One might expect that people share news on social media because they believe it is true. Accordingly, the widespread sharing of false content is often taken as evidence of widespread false beliefs [62,63]. However, recent work has shown that social media sharing judgments can actually be quite divergent from judgments about accuracy [21,44]. For example, participants who were asked about the accuracy of a set of headlines rated true headlines as much more accurate than false headlines; but, when asked whether they would share the headlines, veracity had little impact on sharing intentions – both in the context of political headlines [21] (Figure 2A) and headlines about COVID-19 [44]. As a result, sharing intentions for false headlines were much higher than assessments of their truth (e.g., 91% higher in Figure 2A), indicating that many people were apparently willing to share content that they could have identified as being inaccurate.

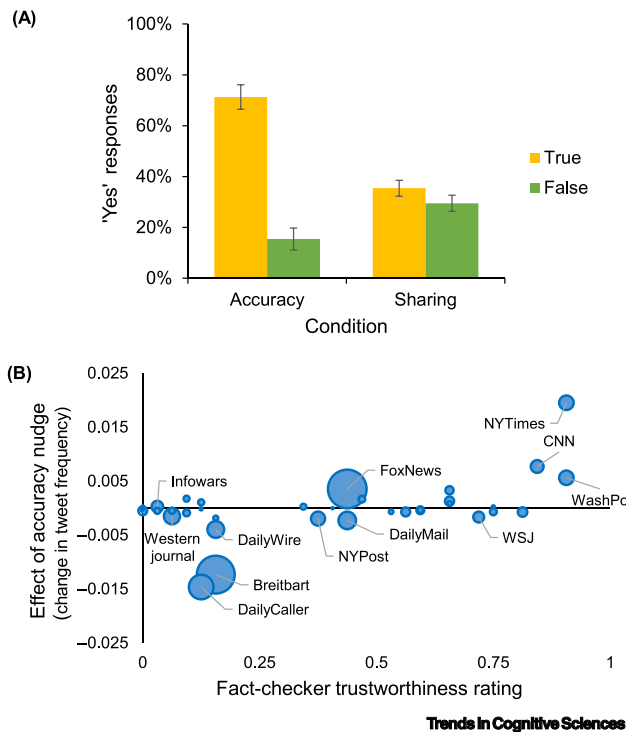


Figure 2. Social Media Sharing Does Not Necessarily Imply Belief, and Accuracy Prompts Improve the Quality of Content That People Share. (A) Data from study 1 in [21] where US participants ($n = 1002$) from Amazon Mechanical Turk (MTurk) judged political headlines. Participants were presented with a set of headlines and were either asked to indicate whether they thought the headlines were accurate (accuracy condition) or whether they would consider sharing them on social media (sharing condition). Although participants were much more likely to rate true headlines versus false headlines as being accurate in the accuracy condition, headline veracity had little impact on sharing intentions in the sharing condition (i.e., there was an interaction between condition and veracity, $b = -0.50$, $t(18089) = -15.37$, $P = <0.001$). Error bars indicate 95% confidence intervals. (B) Data from study 7 in [21] where Twitter users ($n = 5379$) who had recently shared links to websites that regularly produce misleading and hyperpartisan content were sent an unsolicited message asking

them to rate the accuracy of a single nonpolitical headline. The news sites to which users retweeted links in the 24 h after receiving the message were compared with the links retweeted by participants who had not yet received the message (and the date of message delivery was randomized to allow causal inference). The x axis indicates the trust score given to each outlet by eight professional fact-checkers. The y axis indicates the fraction of rated links to each outlet in the 24 h after the intervention minus the fraction of links to each outlet among not-yet-messaged users. The size of each dot is proportional to the number of premessaging posts with links to that outlet. Domains with >500 premessaging posts are labeled. As can be seen, the message significantly improved the quality of new sources shared, with the change in relative tweet frequency being strongly related to fact-checker ratings (domains weighted by number of pretreatment posts; $r(52) = 0.74$, $P = <0.001$). News sources: Breitbart; CNN, Cable News Network; DailyCaller; DailyMail, *The Daily Mail*; DailyWire; FoxNews, Fox News Channel; Infowars; NYPPost, *The New York Post*; NYTimes, *The New York Times*; WashPo, *The Washington Post*; Western Journal; WSJ, *The Wall Street Journal*.

To shed light on this disconnect between accuracy judgments and sharing intentions, a recent study examined the impact of asking participants to rate the perceived accuracy of each headline immediately before deciding whether they would be willing to share it on social media [21]. This experiment helps to distinguish between three distinct explanations for the accuracy–sharing dissociation. The confusion-based account posits that people genuinely (but mistakenly) believe that the false claims they share are probably true. Consistent with this proposal, of the false headlines that were shared in the baseline condition of [21], 33% were both believed and shared when participants were asked directly about accuracy – however, this leaves the remaining 67% of sharing unexplained by confusion.

The preference-based account is rooted in the idea that people place their preference for political identity (or related motives such as virtue signaling [64]) above the truth, and thus share politically consistent false content on social media despite recognizing that it is probably not true. This purposeful sharing could be motivated, for example, by an effort to further one's political agenda [18], to sow chaos [65], or to share news that would be interesting if it turned

out to be true [66]. Of the false headlines that were shared in the baseline condition of [21], 16% of the headlines were shared despite being identified as inaccurate. Thus, although purposeful sharing occurs, it seems unlikely to explain the bulk of false or misleading content that is shared online.

Finally – and most consistent with the earlier focus on a lack of reflective thinking being a source of misjudgments – the inattention-based account argues that people have a strong preference to only share accurate content, but that the social media context distracts them from this preference. Consistent with this account, asking participants to rate the accuracy of each headline before deciding whether to share it decreased sharing of false headlines by 51% relative to the baseline condition [21] – suggesting that inattention to accuracy was responsible for roughly half of the misinformation sharing in the experiment.

In a similar vein, work on social media behavior often emphasizes the importance of the 'attention economy' where factors relating to engagement (likes, shares, comments, clicks, etc.) are selected for in social media environments [8,60,67–69]. Accordingly, sharing of low-quality news content on Facebook is associated with ideological extremity [70] and ideological concordance is a much stronger predictor of sharing than it is of belief [21]. Furthermore, analytic thinking is not only associated with more truth discernment, as described earlier, but is also associated with having more discerning sharing intentions in survey experiments (both in terms of false [35,44] and hyperpartisan content [35]) and with the actual sharing of more reliable news outlets on Twitter [71]. One striking possibility is that the social media context itself distracts people from prioritizing the truth when they decide what to share – and, in some cases, may actively promote antisocial behavior [72] and hamper the exercise of analytic and critical thinking. Social media may be both a cause and a consequence of increased political engagement, both good and bad [73].

What Can Be Done? Interventions To Fight Fake News

We now turn to the implications of these findings for interventions intended to decrease the spread and impact of online misinformation.

Current Approaches for Fighting Misinformation

As social media companies are, first and foremost, technology companies, a common approach is the automated detection of problematic news via machine learning, natural language processing, and network analysis [74–76]. Content classified as problematic is then down-ranked by the ranking **algorithm** such that users are less likely to see it. However, creating an effective misinformation classifier faces two fundamental challenges. First, truth is not a black-and-white, clearly defined property: even professional fact-checkers often disagree on how exactly to classify content [77,78]. Thus, it is difficult to decide what content and features should be included in training sets, and artificial intelligence approaches run the risk of false positives and, therefore, of unjustified censorship [79]. Second, there is the problem of nonstationarity: misinformation content tends to evolve rapidly, and therefore the features which are effective at identifying misinformation today may not be effective tomorrow. Consider, for example, the rise of COVID-19 misinformation in 2020 – classifiers trained to detect largely political content were likely unequipped to be effective for novel false and misleading claims relating to health.

Another commonly used approach involves attaching warnings to content that professional fact-checkers have found to be false (reviewed in [80,81]). A great deal of evidence indicates that corrections and warnings do successfully reduce misperceptions

[41,81–83] and sharing [49,84,85]. Despite some early evidence that correction checking could backfire and increase belief in false content [86], recent work has shown that these backfire effects are extremely uncommon and are not a cause for serious concern [87,88].

There are, however, other reasons to be cautious about the sufficiency of professional fact-checking. Most importantly, fact-checking is simply not scalable – it typically requires substantial time and effort to investigate whether a particular claim is false or misleading. Thus, many (if not most) false claims never get fact-checked. Even for those claims that do eventually get flagged, the process is often slow, such that warnings are likely to be absent during the claim's period of peak viral spreading. Furthermore, warnings are typically only attached to blatantly false news, and not to extremely misleading or biased coverage of events that actually occurred. In addition to straightforwardly undermining the reach of fact-checks, this sparse application of warnings could lead to an 'implied truth' effect where users may assume that (false or misleading) headlines without warnings have actually been verified [84]. Fact-checks often also fail to reach their intended audience [89], and may fade over time [90], provide incomplete protection against familiarity effects [49], and cause corrected users to subsequently share more low-quality and partisan content [91].

Another potential approach that is commonly referenced is emphasizing the publishers of news articles, seeking to leverage the reliance on source cues described earlier. This, in theory, could be effective because people (at least in the USA) are actually fairly good at distinguishing between low- and high-quality publishers [92]. However, experimental evidence on emphasizing news publishers is not very encouraging: Numerous studies find that making source information more salient (or removing it entirely) has little impact on whether people judge headlines to be accurate or inaccurate [37,93–97] (although see [98,99]).

New Approaches for Fighting Misinformation

One potentially promising alternative class of interventions involve a more proactive 'inoculation' or 'prebunking' against misinformation [8,100]. For example, the 'Bad News Game' uses a 10–20 minute interactive tutorial to teach people how to identify fake news in an engaging way [101]. An important limitation of such approaches is that they are 'opt in' – that is, people have to actively choose to engage with the inoculation technique (often for a fairly substantial amount of time – at least in terms of the internet attention span [102]). This is particularly problematic given that those most in need of 'inoculation' against misinformation (e.g., people who are low on cognitive reflection) may be the least likely to seek out and participate in lengthy inoculations. Lighter-touch forms of inoculation that simply present people with information that helps them to identify misinformation (e.g., in the context of climate change [103]) may be more scalable. For example, presenting a simple list of 12 digital media literacy tips improved people's capacity to discern between true and false news in the USA and India [104].

Both fact-checking and inoculation approaches are fundamentally directed toward improving people's underlying knowledge or skills. However, as noted earlier, recent evidence indicates that misinformation may spread on social media not only because people are confused or lack the competency to recognize fake news, but also (or even mostly) because people fail to consider accuracy at all when they make choices about what to share online [21,44]. In addition, as mentioned, people who are more intuitive tend to be worse at distinguishing between true and false news content, both in terms of belief (Figure 1A) and sharing [35,71]. This work suggests that interventions aimed at getting people to slow down and


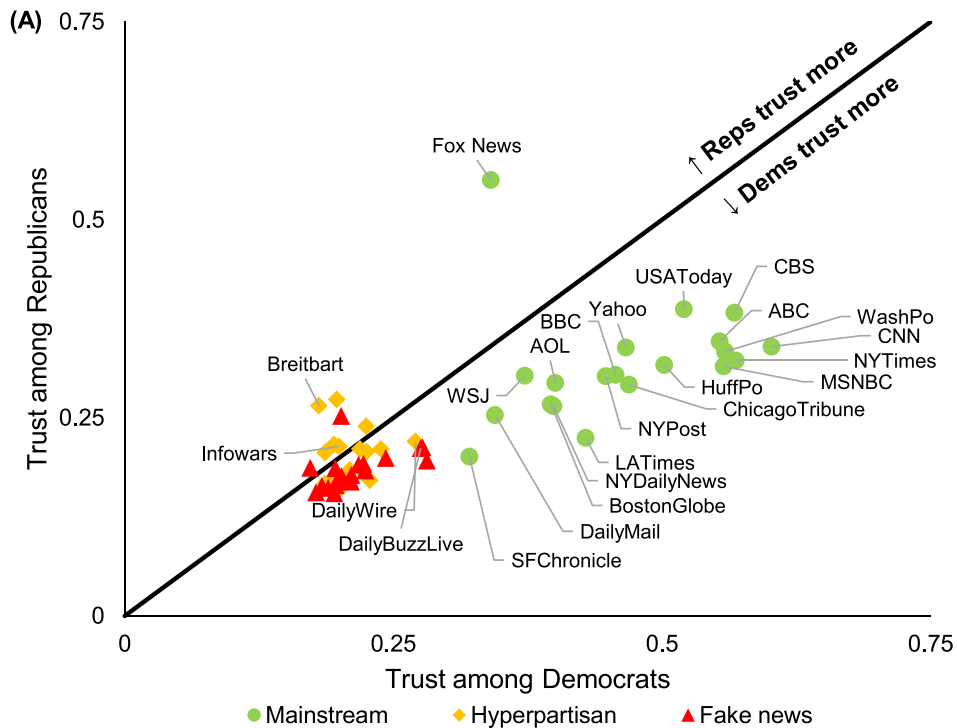
reflect about the accuracy of what they see on social media may be effective in slowing the spread of misinformation.

Indeed, recent research shows that a simple accuracy prompt – specifically, having participants rate the accuracy of a single politically neutral headline (ostensibly as part of a pre-test) before making judgments about social media sharing – improves the extent to which people discern between true and false news content when deciding what to share online in survey experiments [21,44]. This approach has also been successfully deployed in a large-scale field experiment on Twitter, in which messages asking users to rate the accuracy of a politically neutral news headline were sent to thousands of accounts who recently shared links to misinformation sites [21]. This subtle prompt significantly increased the quality of the new they subsequently shared (Figure 2B). Furthermore, survey experiments have shown that asking participants to explain how they know whether a headline is true or false before sharing it increases sharing discernment [105], and having participants rate accuracy at the time of encoding protects against familiarity effects [106]. Relatedly, metacognitive prompts – probing questions that make people reflect – increases resistance to inaccurate information [107].

A major advantage of such accuracy prompts is that they are readily scalable. There are many ways that social media companies, or other interested parties such as governments or civil society organizations, could shift people's attention to accuracy (e.g., through ads, by asking about the accuracy of content that is shared, or via public service announcements, etc.). In addition to scalability, accuracy prompts also have the normative advantage of not relying on a centralized arbiter to determine truth versus falsehood. Instead, they leverage users' own (often latent) ability to make such determinations themselves, preserving user autonomy. Naturally, this will not be effective for everyone all of the time, but it could have a positive effective in the aggregate as one of the various tools used to combat misinformation.

Finally, platforms could also harness the power of human reasoning and the 'wisdom of crowds' to improve the performance of machine-learning approaches. While professional fact-checking is not easily scalable, it is much more tractable for platforms to have large numbers of non-experts rate news content. Despite potential concerns about political bias or lack of knowledge, recent work has found high agreement between layperson crowds and fact-checkers when evaluating the trustworthiness of news publishers: the average Democrat, Republican, and fact-checker all gave fake news and hyperpartisan sites very low trust ratings [92] (Figure 3A). This remained true even when layperson raters were told that their responses would influence social media ranking algorithms, creating an incentive to 'game the system' [108]. However, these studies also revealed a weakness of publisher-based crowd ratings: familiarity with a publisher was necessary (although not sufficient) for trust, meaning that new or niche publishers are unfairly punished by such a rating scheme. One solution to this problem is to have laypeople rate the accuracy of individual articles or headlines (rather than publishers), and to then aggregate these item-level ratings to create average scores for each publisher (Figure 3B). Furthermore, the layperson ratings of the articles themselves are also useful. Analyzing a set of headlines flagged for fact-checking by an internal Facebook algorithm found that the average layperson accuracy rating for fairly small crowds correlated equally well with that of professional fact-checkers as the fact-checkers correlated with each other [77]. Thus, using crowdsourcing to add a 'human in the loop' element to misinformation detection algorithms is promising.

These observations about the utility of layperson ratings have a strong synergy with the aforementioned idea of prompts that shift users' attention to accuracy: periodically asking social



Trends in Cognitive Sciences

(See figure legend at the bottom of the next page.)

media users to rate the accuracy of random headlines both (i) shifts attention to accuracy and thus induces the users to be more discerning in their subsequent sharing, and (ii) generates useful ratings to help inform ranking algorithms.

Concluding Remarks

The spread of misinformation online presents both a scientific puzzle and a practical challenge. The research we have synthesized here shows that the common narrative, whereby failing to differentiate false or misleading news from truth is a symptom of political polarization in a 'post-truth' world, is not an appropriate characterization. Although people do preferentially believe news that aligns with their politics, this occurs as much or more for true headlines compared with false headlines – and thus people are actually more accurate, not less, when judging headlines that are politically concordant. Rather than being bamboozled by partisanship, people often fail to discern truth from fiction because they fail to stop and reflect about the accuracy of what they see on social media. Accordingly, simple prompts that shift people's attention to accuracy increase the quality of news that people share on social media. Approaches of this nature, including providing digital literacy tips, are not hindered by the same issues of scalability related to strict fact-checking approaches – and, in fact, can be combined with crowd-sourced fact-checking to maximize efficiency. Human reasoning, when applied appropriately, can be a powerful salve against the lure of misinformation. Nonetheless, numerous important questions remain unanswered (see Outstanding Questions).

There is also much to be gained for psychology and cognitive science *per se* from investigating online misinformation. Fake news and other forms of online misinformation represent a powerful (and ecologically valid) testing ground for evaluating theories from cognitive, social, and political psychology. More so than many other domains, when studying fake news and misinformation it is often possible to combine traditional laboratory experiments with large-scale social media data [71,109] – or even to conduct actual field experiments on-platform [21,110,111]. Furthermore, these topics also motivate theory development by highlighting a new class of problems in need of further explanation: only a small fraction of online content draws sufficient attention and interest to be shared on social media. Understanding the psychology that underpins the dynamics of social media sharing is an important emerging subfield in psychology [21,67,112] (in addition to other areas such as computer science [76], political science [113], communication [114], and public health [115], among others [116]), driven in large part by concerns about misinformation. Tackling these questions

Figure 3. The Wisdom of Crowds Can Help To Illuminate Information Quality. (A) Data from study 2 of [92] in which US participants ($n = 970$) from Lucid indicated their level of trust in 60 news publishers. Average trust ratings for each source among Democrats (Dems; x axis) and Republicans (Reps; y axis) are shown. Although some partisan differences are evident (e.g., Republicans trust Fox News more than any other publisher), all hyperpartisan and fake news sites received low average trust ratings from members of both parties. Thus, when creating a politically balanced layperson rating by averaging trust from Democrats and Republicans, every mainstream outlet scores higher than every fake news or hyperpartisan outlet; in addition, the layperson ratings were highly correlated with ratings of professional fact-checkers: $r(58) = 0.90, P = <0.001$. (B) Data from study 4 of [93] in which US participants ($n = 1008$) from Amazon Mechanical Turk (MTurk) provided accuracy ratings for the headlines of the 10 most-popular articles from a range of fake news, hyperpartisan, and mainstream accounts. Shown is a layperson outlet-quality score, constructed by averaging the accuracy ratings for the 10 headlines from each outlet, on the y axis, plotted against the rating assigned to each outlet by NewsGuard, a company that hires professional journalists to conduct detailed investigations of the quality of news outlets, on the x axis. The layperson ratings are strongly correlated with the professional journalist ratings: $r(44) = 0.79, P = <0.001$. News source abbreviations are given in Figure 2 legend, additional abbreviations: ABC, American Broadcasting Company; AOL, Aol. (formerly America Online); BBC, British Broadcasting Corporation; bb4sb, BB4SP/Barracuda Brigade; BosGlobe, *The Boston Globe*; CBS, Columbia Broadcasting System (now CBS); ChiTrib, *The Chicago Tribune*; HuffPo, HuffPost (formerly Huffington Post); IJR, Independent Journal Review; LATimes, *The Los Angeles Times*; MSNBC, from Microsoft plus National Broadcasting Company; SFChronicle, *The San Francisco Chronicle*.

Outstanding Questions

What are the impacts of exposure to misinformation? How does exposure affect belief in the specific claims made in the misinformation, more general attitudes (e.g., support for political candidates or policies), and relevant behaviors (e.g., vote choice or health behaviors)? Do these effects differ for misinformation versus reliable news? How much misinformation must be encountered to shift attitudes, beliefs, and behaviors?

Does the design of social media platforms actively promote the spread of misinformation? For example, by inducing distraction and incentivizing or directing attention to factors other than veracity?

How do times of crisis (e.g., the COVID-19 pandemic) affect susceptibility to misinformation and the dynamics of social media behavior?

To what extent do findings about misinformation and social media from the USA and other Western countries generalize to other cultures and other social contexts (e.g., messaging apps such as WhatsApp)?

What are the similarities and differences between the psychology of online misinformation relative to related constellations of beliefs, such as conspiracy ideation, political misperceptions, rumors, and bullshit? How does the psychology of online misinformation compare with misinformation distributed by traditional media such as television and talk radio, or by political elites?

How will findings about current forms of misinformation generalize in the face of the continually evolving nature of online content?

Will social media corporations be willing to invest in, and implement, major changes to fight misinformation, given that their business models are premised on maximizing engagement? How can advocates and policy makers most effectively incentivize them to do so?

How can academics and social media platforms collaborate to study misinformation, given the platforms

will give psychologists the opportunity to demonstrate the power and real-world impact of psychological theory.

near-total control of relevant data, and widespread concerns about privacy?

References

- Lazer, D. et al. (2018) The science of fake news. *Science* 359, 1094–1096
- Wardle, C. (2018) *Information Disorder: The Essential Glossary*, Shorenstein Center on Media, Politics, and Public Policy, Harvard Kennedy School
- Loomba, S. et al. (2021) Measuring the impact of COVID-19 vaccine misinformation on vaccination intent in the UK and USA. *Nat. Hum. Behav.* Published online February 5, 2021. <https://doi.org/10.1038/s41562-021-01056-1>
- Pennycook, G. and Rand, D.G. (2021) Examining false beliefs about voter fraud in the wake of the 2020 Presidential Election. *Harvard Kennedy Sch. Misinformation Rev.* 2, 1–19
- Kaplan, R.L. (2008) Yellow journalism. In *The International Encyclopedia of Communication* (Donsbach, W., ed.), John Wiley & Sons
- Faris, R.M. et al. (2017) *Partisanship, Propaganda, and Disinformation: Online Media and the 2016 U.S. Presidential Election*, Berkman Klein Center for Internet and Society
- Lewandowsky, S. et al. (2020) *Technology and Democracy, Understanding the Influence of Online Technologies on Political Behaviour and Decision-Making*, EU Science Hub
- Kozyreva, A. et al. (2020) Citizens versus the internet: confronting digital challenges with cognitive tools. *Psychol. Sci. Public Interest* 21, 103–156
- Sunstein, C.R. and Vermeule, A. (2009) Conspiracy theories: causes and cures. *J. Polit. Philos.* 17, 202–227
- Lindeman, M. and Aarnio, K. (2007) Superstitious, magical, and paranormal beliefs: An integrative model. *J. Res. Pers.* 41, 731–744
- Berinsky, A.J. (2017) Rumors and health care reform: experiments in political misinformation. *Br. J. Polit. Sci.* 47, 241–262
- Pennycook, G. et al. (2015) On the reception and detection of pseudo-profound bullshit. *Judgm. Decis. Mak.* 10, 549–563
- Amazeen, M.A. et al. (2018) Correcting political and consumer misperceptions: the effectiveness and effects of rating scale versus contextual correction formats. *J. Mass Commun. Q.* 95, 28–48
- Wickens, T. (2002) *Elementary Signal Detection Theory*, Oxford University Press
- Batailler, C. et al. A signal detection approach to understanding the identification of fake news. *Perspect. Psychol. Sci.* (in press)
- Kahan, D.M. (2017) Misconceptions, misinformation, and the logic of identity-protective cognition. *SSRN Electron. J. Cultural Cognition Project Working Paper Series No. 164*, Yale Law School, Public Law Research Paper No. 605, Yale Law & Economics Research Paper No. 575. Published online May 24, 2017. <http://dx.doi.org/10.2139/ssrn.2973067>
- Kahan, D.M. (2013) Ideology, motivated reasoning, and cognitive reflection. *Judgm. Decis. Mak.* 8, 407–424
- Van Bavel, J.J. and Pereira, A. (2018) The partisan brain: an identity-based model of political belief. *Trends Cogn. Sci.* 22, 213–224
- Allcott, H. and Gentzkow, M. (2017) Social media and fake news in the 2016 election. *J. Econ. Perspect.* 31, 211–236
- Pennycook, G. and Rand, D.G. (2019) Lazy, not biased: susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition* 188, 39–50
- Pennycook, G. et al. (2021) Shifting attention to accuracy can reduce misinformation online. *Nature* <http://dx.doi.org/10.1038/s41586-021-03344-2>
- Pereira, A. et al. (2020) Identity concerns drive belief: the impact of partisan identity on the belief and dissemination of true and false news. *PsyArXiv* Published online September 18, 2018. <http://dx.doi.org/10.31234/OSF.IO/7VC5D>
- Vegetti, F. and Mancosu, M. (2020) The impact of political sophistication and motivated reasoning on misinformation. *Polit. Commun.* 37, 678–695
- Faragó, L. et al. (2020) We only believe in news that we doctored ourselves: the connection between partisanship and political fake news. *Soc. Psychol. (Gott)* 51, 77–90
- Drummond, C. et al. (2020) Limited effects of exposure to fake news about climate change. *Environ. Res. Commun.* 2, 081003
- Bago, B. et al. (2020) Fake news, fast and slow: deliberation reduces belief in false (but not true) news headlines. *J. Exp. Psychol. Gen.* 149, 1608–1613
- Tappin, B.M. et al. (2020) Thinking clearly about causal inferences of politically motivated reasoning: why paradigmatic study designs often undermine causal inference. *Curr. Opin. Behav. Sci.* 34, 81–87
- Tappin, B.M. et al. (2020) Rethinking the link between cognitive sophistication and politically motivated reasoning. *J. Exp. Psychol. Gen.* Published online October 29, 2020. <http://dx.doi.org/10.31234/OSF.IO/YUZFJ>
- Tappin, B.M. et al. (2020) Bayesian or biased? Analytic thinking and political belief updating. *Cognition* 204, 1–12
- Baron, J. and Jost, J.T. (2019) False equivalence: are liberals and conservatives in the United States equally biased? *Perspect. Psychol. Sci.* 14, 292–303
- Gerber, A. and Green, D. (1999) Misperceptions about perceptual bias. *Annu. Rev. Polit. Sci.* 2, 189–210
- Leeper, T.J. and Slothuus, R. (2014) Political parties, motivated reasoning, and public opinion formation. *Polit. Psychol.* 35, 129–156
- Friedman, J. (2012) Motivated skepticism or inevitable conviction? Dogmatism and the study of politics. *Crit. Rev.* 24, 131–155
- Bronstein, M.V. et al. (2019) Belief in fake news is associated with delusionality, dogmatism, religious fundamentalism, and reduced analytic thinking. *J. Appl. Res. Mem. Cogn.* 8, 108–117
- Ross, R.M. et al. (2019) Beyond 'fake news': the analytic thinking and the detection of false and hyperpartisan news headlines. *PsyArXiv* Published online November 13, 2019. <http://dx.doi.org/10.31234/osf.io/cgsx6>
- Pehlivanoglu, D. et al. (2020) The role of analytical reasoning and source credibility on the evaluation of real and fake full-length news articles. *PsyArXiv* Published online December 2, 2020. <https://doi.org/10.31234/osf.io/98zwc>
- Pennycook, G. and Rand, D.G. (2019) Who falls for fake news? The roles of bullshit receptivity, overclaiming, familiarity, and analytic thinking. *J. Pers.* 88, 185–200
- Salovich, N.A. et al. (2020) Can confidence help account for and redress the effects of reading inaccurate information? *Mem. Cogn.* 49, 293–310
- Thompson, V.A. et al. (2011) Intuition, reason, and metacognition. *Cogn. Psychol.* 63, 107–140
- Pennycook, G. et al. (2015) What makes us think? A three-stage dual-process model of analytic engagement. *Cogn. Psychol.* 80, 34–72
- Brashier, N.M. et al. (2021) Timing matters when correcting fake news. *Proc. Natl. Acad. Sci. U. S. A.* 118, 1–3
- Amazeen, M.A. and Bucy, E.P. (2019) Confering resistance to digital disinformation: the inoculating influence of procedural news knowledge. *J. Broadcast. Electron. Media* 63, 415–432
- Jones-Jang, S.M. et al. (2019) Does media literacy help identification of fake news? Information literacy helps, but other literacies don't. *Am. Behav.* 65, 371–388
- Pennycook, G. et al. (2020) Fighting COVID-19 misinformation on social media: experimental evidence for a scalable accuracy nudge intervention. *Psychol. Sci.* 31, 770–780
- Kahneman, D. et al. (1982) *Judgments under Uncertainty: Heuristics and Biases*, Cambridge University Press
- Hasher, L. et al. (1977) Frequency and the conference of referential validity. *J. Verbal Learn. Verbal Behav.* 16, 107–112
- Dechens, A. et al. (2010) The truth about the truth: a meta-analytic review of the truth effect. *Personal. Soc. Psychol. Rev.* 14, 238–257
- Allport, F.H. and Lepkin, M. (1945) Wartime rumors of waste and special privilege: why some people believe them. *J. Abnorm. Soc. Psychol.* 40, 3–36
- Pennycook, G. et al. (2018) Prior exposure increases perceived accuracy of fake news. *J. Exp. Psychol. Gen.* 147, 1865–1880

50. Smelter, T.J. and Calvillo, D.P. (2020) Pictures and repeated exposure increase perceived accuracy of news headlines. *Appl. Cogn. Psychol.* 34, 1061–1071
51. Fazio, L.K. et al. (2019) Repetition increases perceived truth equally for plausible and implausible statements. *Psychon. Bull. Rev.* 26, 1705–1710
52. Unkelbach, C. et al. (2019) Truth by repetition: explanations and implications. *Curr. Dir. Psychol. Sci.* 28, 247–253
53. Wang, W.-C. et al. (2016) On known unknowns: fluency and the neural mechanisms of illusory truth. *J. Cogn. Neurosci.* 28, 739–746
54. Cornille, O. et al. (2020) Repetition increases both the perceived truth and fakeness of information: an ecological account. *Cognition* 205, 104470
55. Pompitakpan, C. (2004) The persuasiveness of source credibility: a critical review of five decades' evidence. *J. Appl. Soc. Psychol.* 34, 243–281
56. Zaller, J.R. (1992) *The Nature and Origins of Mass Opinion*, Cambridge University Press
57. Swire, B. et al. (2017) Processing political misinformation: comprehending the Trump phenomenon. *R. Soc. Open Sci.* 4, 160802
58. Avram, M. et al. (2020) Exposure to social engagement metrics increases vulnerability to misinformation. *Harvard Kennedy Sch. Misinformation Rev.* 1, 1–11
59. Quandt, T. (2018) Dark participation. *Media Commun.* 6, 36–48
60. Crockett, M.J. (2017) Moral outrage in the digital age. *Nat. Hum. Behav.* 1, 769–771
61. Martel, C. et al. (2020) Reliance on emotion promotes belief in fake news. *Cogn. Res. Princ. Implic.* 5, 1–20
62. Chatfield, T. (2019) Why we believe fake news. *BBC Future* Published online September 9, 2019. <http://www.bbc.com/future/story/20190905-how-our-brains-get-overloaded-by-the-21st-century>
63. Steinmetz, K. (2018) How your brain tricks you into believing fake news. *Time* Published online August 9, 2019. <https://time.com/5362183/the-real-fake-news-crisis/>
64. Jordan, J.J. and Rand, D.G. (2020) Signaling when no one is watching: a reputation heuristics account of outrage and punishment in one-shot anonymous interactions. *J. Pers. Soc. Psychol.* 118, 57–88
65. Petersen, M.B. et al. (2018) A 'need for chaos' and the sharing of hostile political rumors in advanced democracies. *PsyArXiv* Published online September 1, 2018. <http://dx.doi.org/10.31234/osf.io/6m4ts>
66. Altay, S. et al. (2020) 'If this account is true, it is most enormously wonderful': Interestingness-if-true and the sharing of true and false news. *PsyArXiv* Published online June 4, 2020. <http://dx.doi.org/10.31234/osf.io/tdfh5>
67. Brady, W.J. et al. (2020) The MAD model of moral contagion: the role of motivation, attention and design in the spread of moralized content online. *Perspect. Psychol. Sci.* 15, 978–1010
68. Baek, E.C. et al. (2017) The value of sharing information: a neural account of information transmission. *Psychol. Sci.* 28, 851–861
69. Scholz, C. et al. (2017) A neural model of valuation and information virality. *Proc. Natl. Acad. Sci. U. S. A.* 114, 2881–2886
70. Hopp, T. et al. (2020) Why do people share ideologically extreme, false, and misleading content on social media? A self-report and trace data-based analysis of countermedia content dissemination on Facebook and Twitter. *Hum. Commun. Res.* 46, 357–384
71. Mosleh, M. et al. (2021) Cognitive reflection correlates with behavior on Twitter. *Nat. Commun.* 12, 1–10
72. Efron, D.A. and Raj, M. (2019) Misinformation and morality: encountering fake-news headlines makes them seem less unethical to publish and share. *Psychol. Sci.* 31, 75–87
73. Valenzuela, S. et al. (2019) The paradox of participation versus misinformation: social media, political engagement, and the spread of misinformation. *Digit. J.* 7, 802–823
74. Shu, K. et al. (2017) Fake news detection on social media. *ACM SIGKDD Explor. News.* 19, 22–36
75. Monti, F. et al. (2019) Fake news detection on social media using geometric deep learning. *ArXiv* Published online February 10, 2019. <http://arxiv.org/abs/1902.06673>
76. Zhou, X. et al. (2019) Fake news: fundamental theories, detection strategies and challenges. In *WSDM 2019: Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pp. 836–837. Association for Computing Machinery, Inc
77. Allen, J. et al. (2020) Scaling up fact-checking using the wisdom of crowds. *PsyArXiv* Published online October 2, 2020. <http://dx.doi.org/10.31234/osf.io/9qdzd>
78. Lim, C. (2018) Checking how fact-checkers check. *Res. Polit.* 5 205316801878684
79. Woolley, S. (2020) We're fighting fake news AI bots by using more AI. That's a mistake. *MIT Technol. Rev.* Published online January 8, 2020. <https://www.technologyreview.com/2020/01/08/130983/were-fighting-fake-news-ai-bots-by-using-more-ai-thats-a-mistake/>
80. Walter, N. et al. (2020) Fact-checking: a meta-analysis of what works and for whom. *Polit. Commun.* 37, 350–375
81. Nieminen, S. and Rapeli, L. (2019) Fighting misperceptions and doubting journalists' objectivity: a review of fact-checking literature. *Polit. Stud. Rev.* 17, 296–309
82. Nyhan, B. et al. (2019) Taking fact-checks literally but not seriously? The effects of journalistic fact-checking on factual beliefs and candidate favorability. *Polit. Behav.* 42, 939–960
83. Ecker, U.K.H. et al. (2019) The effectiveness of short-format refutational fact-checks. *Br. J. Psychol.* 111, 36–54
84. Pennycook, G. et al. (2020) The implied truth effect: attaching warnings to a subset of fake news stories increases perceived accuracy of stories without warnings. *Manag. Sci.* 66, 4944–4957
85. Yaqub, W. et al. (2020) Effects of credibility indicators on social media news sharing intent. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (vol. 20)*, pp. 1–14, Association for Computing Machinery, Inc
86. Nyhan, B. and Reifer, J. (2010) When corrections fail: the persistence of political misperceptions. *Polit. Behav.* 32, 303–330
87. Wood, T. and Porter, E. (2019) The elusive backfire effect: mass attitudes' steadfast factual adherence. *Polit. Behav.* 41, 135–163
88. Guess, A.M. and Coppock, A. (2018) Does counter-attitudinal information cause backlash? Results from three large survey experiments. *Br. J. Polit. Sci.* 50, 1497–1511
89. Guess, A.M. et al. (2020) Exposure to untrustworthy websites in the 2016 US election. *Nat. Hum. Behav.* 4, 472–480
90. Swire, B. et al. (2017) The role of familiarity in correcting inaccurate information. *J. Exp. Psychol. Learn. Mem. Cogn.* 43, 1948–1961
91. Mosleh, M. et al. (2021) Perverse consequences of debunking in a Twitter field experiment: being corrected for posting false news increases subsequent sharing of low quality, partisan, and toxic content. In *Proceedings of the 2021 CHI Conference of Human Factors in Computing Systems*, <http://dx.doi.org/10.1145/3411764.3445642>
92. Pennycook, G. and Rand, D.G. (2019) Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proc. Natl. Acad. Sci. U. S. A.* 116, 2521–2526
93. Dias, N. et al. (2020) Emphasizing publishers does not effectively reduce susceptibility to misinformation on social media. *Harvard Kennedy Sch. Misinformation Rev.* 1, 1–12
94. Austin, E.W. and Dong, Q. (1994) Source v. content effects on judgments of news believability. *J. Mass Commun. Q.* 71, 973–983
95. Jakesch, M. et al. (2019) The role of source, headline and expressive responding in political news evaluation. *SSRN Electron.* Published online January 8, 2019. <https://doi.org/10.2139/ssrn.3306403>
96. Shen, C. et al. (2019) Fake images: the effects of source, intermediary, and digital media literacy on contextual assessment of image credibility online. *New Media Soc.* 21, 438–463
97. Tsang, S.J. (2020) Motivated fake news perception: the impact of news sources and policy support on audiences' assessment of news fakeness. *J. Mass Commun. Q.* Published online August 31, 2020. <https://doi.org/10.1177/1077699020952129>
98. Kim, A. et al. (2019) Combating fake news on social media with source ratings: the effects of user and expert reputation ratings. *J. Manag. Inf. Syst.* 36, 931–968
99. Nadarevic, L. et al. (2020) Perceived truth of statements and simulated social media postings: an experimental investigation

- of source credibility, repeated exposure, and presentation format. *Cogn. Res. Princ. Implic.* 5, 56
100. Compton, J. (2013) Inoculation theory. In *The SAGE Handbook of Persuasion: Developments in Theory and Practice* (Dillard, J.P. and Shen, L., eds), pp. 220–236, SAGE Publications, Inc
 101. Roozenbeek, J. and van der Linden, S. (2019) Fake news game confers psychological resistance against online misinformation. *Palgrave Commun.* 5, 65
 102. Lorenz-Spreen, P. et al. (2019) Accelerating dynamics of collective attention. *Nat. Commun.* 10, 1759
 103. Cook, J. et al. (2017) Neutralizing misinformation through inoculation: exposing misleading argumentation techniques reduces their influence. *PLoS One* 12, e0175799
 104. Guess, A.M. et al. (2020) A digital media literacy intervention increases discernment between mainstream and false news in the United States and India. *Proc. Natl. Acad. Sci. U. S. A.* 117, 15536–15545
 105. Fazio, L. (2020) Pausing to consider why a headline is true or false can help reduce the sharing of false news. *Harvard Kennedy Sch. Misinformation Rev.* 1, 1–8
 106. Brashier, N.M. et al. (2020) An initial accuracy focus prevents illusory truth. *Cognition* 194, 104054
 107. Salovich, N.A. and Rapp, D.N. (2020) Misinformed and unaware? Metacognition and the influence of inaccurate information. *J. Exp. Psychol. Learn. Mem. Cogn.* Published online November 5, 2020. <https://psycnet.apa.org/doi/10.1037/xlm0000977>
 108. Epstein, Z. et al. (2020) Will the crowd game the algorithm? Using Layperson Judgments To Combat Misinformation On Social Media By Downranking Distrusted Sources. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*
 109. Carpenter, J. et al. (2018) The impact of actively open-minded thinking on social media communication. *Judgm. Decis. Mak.* 13, 562–574
 110. Munger, K. (2017) Tweetment effects on the Tweeted: experimentally reducing racist harassment. *Polit. Behav.* 39, 629–649
 111. Mosleh, M. et al. (2021) Shared partisanship dramatically increases social tie formation in a Twitter field experiment. *Proc. Natl. Acad. Sci. U. S. A.* 118, e2022761118
 112. Brady, W.J. et al. (2020) Attentional capture helps explain why moral and emotional content go viral. *J. Exp. Psychol. Gen.* 149, 746–756
 113. Tucker, J. et al. (2018) Social media, political polarization, and political disinformation: a review of the scientific literature. *SSRN Electron. J.* Published online March 21, 2018. <https://doi.org/10.2139/ssrn.3144139>
 114. Li, J. (2020) Toward a research agenda on political misinformation and corrective information. *Polit. Commun.* 37, 125–135
 115. Wang, Y. et al. (2019) Systematic literature review on the spread of health-related misinformation on social media. *Soc. Sci. Med.* 240, 112552
 116. Paletz, S.B.F. et al. (2019) *A Multidisciplinary Framework of Information Propagation Online*, Springer
 117. Grinberg, N. et al. (2019) Fake news on twitter during the 2016 U.S. Presidential election. *Science* 363, 374–378
 118. Guess, A.M. et al. (2019) Less than you think: prevalence and predictors of fake news dissemination on Facebook. *Sci. Adv.* 5, eaau4586
 119. Bradshaw, S. et al. (2020) Sourcing and automation of political news and information over social media in the United States, 2016–2018. *Polit. Commun.* 37, 173–193
 120. Rogers, R. (2020) The scale of Facebook’s problem depends upon how ‘fake news’ is classified. *Harvard Kennedy Sch. Misinformation Rev.* 1, 1–15
 121. Halpern, D. et al. (2019) From belief in conspiracy theories to trust in others: which factors influence exposure, believing and sharing fake news. In *Social Computing and Social Media. Design, Human Behavior and Analytics. HCLII 2019. Lecture Notes in Computer Science* (vol 11578) (Meiselwitz, G., ed.), pp. 217–232, Springer, Cham
 122. Zimmermann, F. and Kohring, M. (2020) Mistrust, disinforming news, and vote choice: a panel survey on the origins and consequences of believing disinformation in the 2017 German Parliamentary Election. *Polit. Commun.* 37, 215–237
 123. Ditto, P.H. et al. (2019) At least bias is bipartisan: a meta-analytic comparison of partisan bias in liberals and conservatives. *Perspect. Psychol. Sci.* 14, 273–291
 124. Clark, C.J. and Winegard, B.M. (2020) Tribalism in war and peace: the nature and evolution of ideological epistemology and its significance for modern social science. *Psychol. Inq.* 31, 1–22
 125. Druckman, J.N. and McGrath, M.C. (2019) The evidence for motivated reasoning in climate change preference formation. *Nat. Clim. Chang.* 9, 111–119
 126. Kahan, D.M. et al. (2012) The polarizing impact of science literacy and numeracy on perceived climate change risks. *Nat. Clim. Chang.* 2, 732–735
 127. Evans, J.S.B.T. et al. (1983) On the conflict between logic and belief in syllogistic reasoning. *Mem. Cogn.* 11, 295–306
 128. Bago, B. et al. (2020) Reasoning about climate change. *PsyArXiv* Published online December 13, 2020. <https://doi.org/10.31234/osf.io/vcpkb>
 129. Koehler, J.J. (1993) The influence of prior beliefs on scientific judgments of evidence quality. *Organ. Behav. Hum. Decis. Process.* 56, 28–55
 130. Hahn, U. and Harris, A.J.L. (2014) What Does It Mean to be Biased. Motivated Reasoning and Rationality. In *Psychology of Learning and Motivation – Advances in Research and Theory* (vol. 61), pp. 41–102, Academic Press
 131. McPhetres, J. and Pennycook, G. (2020) Science beliefs, political ideology, and cognitive sophistication. *OSF Preprints* Published online November 19, 2019. <https://doi.org/10.31219/osf.io/ad9v7>
 132. Pennycook, G. et al. (2020) On the belief that beliefs should change according to evidence: Implications for conspiratorial, moral, paranormal, political, religious, and science beliefs. *Judgm. Decis. Mak.* 15, 476–498
 133. Bergan, D.E. (2021) Does ability contribute to partisan bias? Evaluating scientific research about political topics. *Commun. Stud.* Published online February 4, 2021. <https://doi.org/10.1080/10510974.2021.1876124>
 134. De Neys, W. (2017) *Dual Process Theory 2.0*, Routledge
 135. Evans, J.S.B.T. and Stanovich, K.E. (2013) Dual-process theories of higher cognition: Advancing the debate. *Perspect. Psychol. Sci.* 8, 223–241
 136. Frederick, S. (2005) Cognitive reflection and decision making. *J. Econ. Perspect.* 19, 25–42
 137. Toplak, M.E. et al. (2011) The cognitive reflection test as a predictor of performance on heuristics-and-biases tasks. *Mem. Cogn.* 39, 1275–1289
 138. Pennycook, G. et al. (2015) Everyday consequences of analytic thinking. *Curr. Dir. Psychol. Sci.* 24, 425–432
 139. Swami, V. et al. (2014) Analytic thinking reduces belief in conspiracy theories. *Cognition* 133, 572–585
 140. Pennycook, G. and Rand, D.G. (2019) Cognitive reflection and the 2016 U.S. presidential election. *Pers. Soc. Psychol. Bull.* 45, 224–239