

MIT Open Access Articles

Fake news, fast and slow: Deliberation reduces belief in false (but not true) news headlines

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Bago, Bence et al. "Fake news, fast and slow: Deliberation reduces belief in false (but not true) news headlines." *Journal of Experimental Psychology: General*, 129, 8 (August 2020): 1608-1613 © 2020 The Author(s)

As Published: 10.1037/XGE0000729

Publisher: American Psychological Association (APA)

Persistent URL: <https://hdl.handle.net/1721.1/130414>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



**Fake news, fast and slow:
Deliberation reduces belief in false (but not true) news headlines**

In press at Journal of Experimental Psychology: General
DOI: <http://dx.doi.org/10.1037/xge0000729>

Bence Bago^{1*}, David G. Rand^{2,3}, & Gordon Pennycook⁴

¹Institute for Advanced Study in Toulouse, University of Toulouse Capitole

²Sloan School, Massachusetts Institute of Technology

³Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology

⁴Hill/Levene Schools of Business, University of Regina

*Corresponding author: Bence Bago, Institute for Advanced Study in Toulouse, University of Toulouse Capitole, 21 Allée de Brienne, 31015 Toulouse, France, e-mail: bencebagok@gmail.com

Abstract

What role does deliberation play in susceptibility to political misinformation and “fake news”? The “Motivated System 2 Reasoning” account posits that deliberation causes people to fall for fake news because reasoning facilitates identity-protective cognition and is therefore used to rationalize content that is consistent with one’s political ideology. The classical account of reasoning instead posits that people ineffectively discern between true and false news headlines when they *fail* to deliberate (and instead rely on intuition). To distinguish between these competing accounts, we investigated the causal effect of reasoning on media truth discernment using a two-response paradigm. Participants (N= 1635 MTurkers) were presented with a series of headlines. For each, they were first asked to give an initial, intuitive response under time pressure and concurrent working memory load. They were then given an opportunity to re-think their response with no constraints, thereby permitting more deliberation. We also compared these responses to a (deliberative) one-response baseline condition where participants made a single choice with no constraints. Consistent with the classical account, we found that deliberation corrected intuitive mistakes: subjects believed false headlines (but not true headlines) more in initial responses than in either final responses or the unconstrained 1-response baseline. In contrast – and inconsistent with the Motivated System 2 Reasoning account – we found that political polarization was equivalent across responses. Our data suggest that, in the context of fake news, deliberation facilitates accurate belief formation and not partisan bias.

Keywords: fake news, misinformation, dual-process theory, two-response paradigm

Although inaccuracy in news is nothing new, so-called “fake news” – “fabricated information that mimics news media content in form but not in organisational process or intent” (Lazer et al., 2018, p. 1094) – has become a focus of attention in recent years. Fake news represents an important test-case for psychologists: What is it about human reasoning that allows people to fall for blatantly false content?

Here we consider this question from a dual-process perspective, which distinguishes between intuitive versus deliberative cognitive processing (Evans & Stanovich, 2013; Kahneman, 2011). The theory posits that intuition allows for quick automatic responses that are often based on heuristic cues, while effortful deliberation can override and correct intuitive responses.

With respect to misinformation and the formation of (in)accurate beliefs, there is substantial debate about the roles of intuitive versus deliberative processes. In particular, there are two major views: the “Motivated System 2 Reasoning” (MS2R) account and the classical reasoning account. According to the MS2R account, people engage in deliberation to protect their (often political) identities and to defend their pre-existing beliefs. As a result, deliberation *increases* partisan bias (Charness & Dave, 2017; Kahan, 2013, 2017; Kahan et al., 2012; Sloman & Rabb, 2019).¹ In the context of evaluating news, this means that increased deliberation will lead to increased political polarization and decreased ability to discern true from false. Support for this account comes from studies that correlate deliberativeness with polarization. For example, highly numerate people are more likely to be polarized on a number of political issues, including climate change (Kahan et al., 2012) and gun control (Kahan, Peters, Dawson, & Slovic, 2017). Furthermore, Kahan, Peters, Dawson,

¹ Various accounts of motivated reasoning other than MS2R have been offered (e.g., Dawson, Gilovich, & Regan, 2002; Haidt, 2001; Mercier & Sperber, 2011), and these accounts may make differing predictions regarding the role of deliberation. Here we focus on the MS2R account (Kahan, 2017), as this account makes clear and specific predictions about the connection between greater deliberation and increased partisan bias.

and Slovic (2017) experimentally manipulated the political congruence of information they presented to participants and found that the ratings of highly numerate participants responded more to the congruence manipulation.

The classical account of reasoning, in contrast, argues that when people engage in deliberation, it typically helps uncover the truth (Evans, 2010; Evans & Stanovich, 2013; Pennycook & Rand, 2019a; Shtulman & McCallum, 2014; Stanovich, 2011; Swami, Voracek, Stieger, Tran, & Furnham, 2014). In the context of misinformation, the classical account therefore posits that it is the lack of deliberation that promotes belief in fake news, while deliberation results in greater truth discernment (Pennycook & Rand, 2019a). Support for the classical account comes from correlational evidence that people who are dispositionally more deliberative are better able to discern between true and false news headlines, regardless of the ideological alignment of the content (Pennycook and Rand 2019a; see also Bronstein, Pennycook, Bear, Rand, & Cannon, 2019; Pennycook & Rand, 2019b). Relatedly, it has been shown that people update their prior beliefs when presented with evidence about the scientific consensus regarding anthropogenic climate change, regardless of their prior motivation or political orientation (Van der Linden, Leiserowitz, and Maibach (2018); see also Lewandowsky, Gignac, & Vaughan, 2013); and that training to detect fake news decreases belief regardless of partisanship (Roozenbeek & van der Linden, 2019a, 2019b). Although not directly manipulating deliberation, these results suggest that engaging in reasoning leads to more accurate, rather than more polarized, beliefs.

To differentiate between these motivated and classical accounts, the key question, then, is this: When assessing news, does deliberation cause an increase in polarization or in accuracy?

Here we shed new light on this question by experimentally investigating the causal link between deliberation and polarization (MS2R) versus correction (classical reasoning).

Specifically, we used the two-response paradigm, in which participants are presented with the same news headline twice. First, they are asked to give a very quick, intuitive response under time pressure and working memory load (Bago & De Neys, 2019). After this, they are presented with the task again and asked to give a final response without time pressure or working memory load (thus allowing unrestricted deliberation). This paradigm has been shown to reliably manipulate the relative roles of intuition and deliberation across a range of tasks (e.g., Bago & De Neys, 2017, 2019; Thompson, Turner, & Pennycook, 2011).

The classical account predicts that false headlines – but not true headlines – should be judged to be less accurate in deliberative (final) responses compared to intuitive (initial) responses; and that this should be the case regardless of whether the headlines are politically concordant (e.g., a headline with a Pro-Democratic lean for a Democrat) or discordant (e.g., a headline with a Pro-Democratic lean for a Republican). In contrast, the MS2R account predicts that politically discordant headlines should be judged to be less accurate – and politically concordant headlines judged to be more accurate – for deliberative responses compared to intuitive responses, regardless of whether the headlines are true or false.

Method

Data, preregistrations of sample sizes and primary analyses, and supplementary materials are available on the Open Science Framework: <https://osf.io/egy8p>.

The preregistered sample for this study was 1000 online participants recruited from Mechanical Turk (Horton, Rand, & Zeckhauser, 2011); 600 for the two-response experiment, and 400 for a one-response baseline condition (participants in previous experiments of ours on this topic were not allowed to participate). In total, 1012 participants were recruited (503 females, $M_{\text{age}} = 36.9$ years). The research project was approved by the University of Regina and MIT Research Ethics Boards.

Participants rated the accuracy of 16 actual headlines taken from social media; four each of Republican-consistent false, Republican-consistent true, Democrat-consistent false, and Democrat-consistent true. Headlines were presented in a random order and randomly sampled from a pool of 24 total headlines (from Pennycook and Rand, 2019a). For each headline, participants were asked “Do you think this headline describes an event that actually happened in an accurate way?” with the response options “Yes/No” (the order of Yes/No v. No/Yes was counterbalanced across participants).

In the one-response baseline condition, participants merely rated the 16 headlines, taking as long as they desired for each. In the two-response experiment, participants made an initial response in which the extent of deliberation was minimized by having participants complete a load task (memorizing a pattern of five dots in a 4X4 matrix, see Bago & De Neys, 2019) and respond within 7 seconds (the average reading time in an N = 104 pre-test). They were then presented with the same headline again – with no time deadline or load – and asked to give a final response.

After rating the 16 headlines, participants completed a variety of demographics, including the Cognitive Reflection Test (Frederick, 2005; Thomson & Oppenheimer, 2016; CRT) and a measure of support for the Republican versus Democratic Party (which we used to classify headlines as politically concordant versus discordant). See Supplementary Materials for further methodological details.

We analyze the results using mixed-effect logistic regression models, with headlines and subjects as random intercepts. Any analysis that is not preregistered is labeled as *post hoc*. We necessarily excluded the 4.1% of trials where individuals missed the initial response deadline. We also preregistered that we would exclude trials where individuals gave an incorrect response to the load task, but we found a significant correlation between score on the CRT and performance on the cognitive load task, $r = 0.11$, $p < 0.0001$. Thus, we kept the

incorrectly solved load trials to avoid a possible selection bias. Note that, for completeness, we also ran the analysis with the preregistered exclusions and there are no notable deviations from the results presented here (see Supplementary Materials). Furthermore, 14 participants did not give a response to our political ideology question and were also excluded from subsequent analyses. As preregistered, we excluded no trials when comparing the one-response baseline to the final response of the two-response paradigm to avoid selection bias (apart from the 14 participants who did not answer the ideology question).

Results

Politically neutral pre-test. We begin by reporting the results of a pre-test that used politically neutral headlines ($N = 623$; see Supplementary Materials for details). As there is no motivation to (dis)believe these headlines, the straightforward prediction is that deliberation would reduce the perceived accuracy of false (but not true) headlines. Indeed, in the two-response experiment there was a significant interaction between headline veracity and response number (initial vs final), $b = 0.47$, 95% CI = [0.29, 0.65], $p < 0.0001$. Similarly, when comparing across conditions, there was a significant interaction between headline veracity and condition (one-response baseline vs two-response experiment), using either the initial response, $b = 0.61$, 95% CI = [0.42, 0.79], $p < 0.0001$, or final response, $b = 0.23$, 95%

CI = [0.04, 0.41], $p = 0.018$, from the two-response experiment (see Figure 1). Deliberation increased ability to discern true versus false politically neutral headlines.




Figure 1. Percentage of true versus false politically neutral headlines that subjects rated as accurate across conditions. Error bars are 95% CIs.

Within-subject analysis. We now turn to our main experiment, where subjects judged political headlines, to adjudicate between the MS2R and classical accounts (see Figure 2). First, we compare initial (intuitive) versus final (deliberative) responses within the two-response experiment to investigate the causal effect of deliberation within-subject. Consistent with the classical account, we found a significant interaction between headline veracity and response number, $b = 0.36$, 95% CI = [0.2, 0.52], $p < 0.0001$, such that final responses rated false (but not true) news as less accurate relative to initial answers. Moreover, inconsistent with the MS2R account, there was no interaction between political concordance and response number, $b = 0.004$, 95% CI = [-0.16, 0.17], $p = 0.96$, and no three-way interaction between response type, political concordance, and headline veracity, $b = 0.03$, 95% CI = [-0.14, 0.21] $p = 0.72$. Thus, people were more likely to correct their response after deliberation, regardless

of whether the item was concordant or discordant with their political beliefs. Naturally, concordance had *some* effect – people rated politically concordant headlines as more accurate than discordant ones, $b = -0.21$, 95% CI = [-0.34, -0.07], $p = 0.003$ – but this was equally true for initial and final responses.

There was also a significant interaction between political concordance and headline veracity, $b = -0.3$, 95% CI = [-0.47, -0.14], $p = 0.0003$, such that the difference between politically concordant and discordant news was larger for real items than for fake items – that is, people were *more* politically polarized for real news than for fake news – but, again, this was equally true for initial versus final responses. Finally, we also found significant main effects of veracity (perceived accuracy was lower for false than true news), $b = 1.56$, 95% CI = [1.14, 1.98] $p < 0.0001$, and response type (perceived accuracy was lower for final than initial responses), $b = -0.38$, 95% CI = [-0.52, -0.25], $p < 0.0001$.

We also examined the role of dispositional differences in deliberativeness (as measured by performance on the CRT). As described in detail in the Supplementary Materials, we replicated prior findings that people who scored higher on the CRT were better at discerning true versus false headlines, and we found significant interactions with response number such that this relationship between CRT and discernment was stronger for final responses than initial responses (although still present for initial responses).




Figure 2. Percentage of true versus false political headlines that subjects rated as accurate across conditions and political concordance. Error bars are 95% CIs.

Between-subject analysis. Finally, we compare perceived accuracy ratings in the two-response experiment with ratings from the one-response baseline (see Figure 2). We first report a *post hoc* analysis comparing the initial (intuitive) response from the two-response experiment with the one-response baseline. This recapitulates a standard load/time pressure experiment, in which some subjects respond under load/pressure while others do not. We found a significant interaction between headline veracity and condition; $b = 0.34$, 95% CI = [0.17, 0.51], $p < 0.0001$; concordance and veracity; $b = -0.31$, 95% CI = [-0.48, -0.15], $p = 0.0002$; and veracity, condition, and concordance, $b = 0.21$, 95% CI = [0.03, 0.39], $p = 0.035$. Load/time pressure increased perceived accuracy of fake headlines regardless of political concordance. Load/time pressure had no effect for politically concordant real headlines, but did decrease perceived accuracy of politically discordant real headlines. Therefore, deliberation causes an increase in truth discernment for both discordant and concordant headlines.

We conclude by comparing the final (deliberative) response from the two-response experiment with the one-response baseline. This allows us to test whether forcing subjects to

report an initial response in the two-response experiment had some carryover (e.g. anchoring) effect on their final response. Although there was no significant interaction between veracity and condition, $b = 0.03$, 95% CI = [-0.14, 0.2] $p = 0.74$, there was a significant interaction between veracity and concordance $b = -0.28$, 95% CI = [-0.44, -0.11], $p = 0.0009$, and a significant three-way interaction between veracity, condition, and concordance, $b = 0.19$, 95% CI = [0.01, 0.37], $p = 0.037$. Politically discordant items showed an anchoring effect whereby perceived accuracy of fake headlines was lower – and perceived accuracy of real headlines was higher – for the one-response baseline relative to the final response of the two-response condition. For politically concordant items, however, there was no such anchoring effect. Together with the significant anchoring effect among politically neutral headlines observed in our pre-test, this suggests that there is something unique about politically concordant items when it comes to anchoring.

Discussion

What is the role of deliberation in assessing the truth of news? We found experimental evidence supporting the classical account over the MS2R account. Broadly, we found that people made fewer mistakes in judging the veracity of headlines – and in particular were less likely to believe false claims – when they deliberated, regardless of whether or not the headlines aligned with their ideology. Conversely, we found no evidence that deliberation influenced the level of partisan bias/polarization.

Theoretical implications

These observations have important implications for both theory and practice. From a theoretical perspective, our results provide the first causal evidence regarding the “corrective” role of deliberation in media truth discernment. There has been a spirited debate regarding the role of deliberation and reasoning among those studying misinformation and political thought, but this debate has proceeded without causal evidence regarding the impact of

manipulating deliberation on polarization versus correction. To our knowledge, our experiment is the first that enables us to do so – and provides clear support for the classical account of reasoning.

Limitations and future directions

Using similar methods to test the role of deliberation in the continued influence effect (CIE, Johnson & Seifert, 1994), wherein people continue to believe in misinformation even after it was retracted or corrected (Lewandowsky, Ecker, Seifert, Schwarz, & Cook, 2012), is a promising direction for future work. So too is examining the impact of deliberation on the many (psychological) factors which have been shown to influence the acceptance of corrections, such as trust in the source of original information (Swire, Berinsky, Lewandowsky, & Ecker, 2017), underlying worldview or political orientation (Ecker & Ang, 2019), and strength of encoding of the information (Ecker, Lewandowsky, Swire, & Chang, 2011). For example, deliberation might make it easier to accept corrections and update beliefs. Relatedly, the computations taking place during deliberation are underspecified, and therefore future work could benefit from developing formalized, computational models that better characterize underlying computations, such as the Decision by Sampling model (Stewart, Chater, & Brown, 2006).

One limitation of the current work is that it was conducted on non-nationally representative samples from MTurk. However, it is not imperative for us to have an ideologically balanced/representative sample, as we are not making comparisons between holders of one ideology versus another. Instead, we investigate motivated reasoning – which should apply to both Democrats and Republicans – by comparing discordant versus concordant headlines (collapsing across Democrats and Republicans). It would be interesting for future work to replicate our results using a more representative sample to investigate the potential for partisan asymmetries in the impact of deliberation.

Another potential concern is that our sample may not contain the people who are the most susceptible to misinformation, given that the baseline levels of belief in fake news we observe are low (Kahan, 2018). This problem is endemic in survey-based research on misinformation. Future work could potentially address such issues by using advertising on social media to recruit participants who have actually shared misinformation in the past.

Practical implications

From a practical perspective, the proliferation of false headlines has been argued to pose potential threats to democratic institutions and people, by increasing apathy and polarization or even inducing violent behavior (Lazer et al., 2018). Thus, there is a great deal of interest around developing policies to combat the influence of misinformation. Such policies should be grounded in an understanding of the underlying psychological processes that lead people to fall for inaccurate content. Our results suggest that fast, intuitive (likely emotional; Martel, Pennycook, & Rand, 2019) processing plays an important role in promoting belief in false content – and therefore that interventions that promote deliberation may be effective. Relatedly, this suggests that the success of fake news on social media may be related to users’ tendency to scroll quickly through their newsfeeds, and the use of highly emotionally engaging content by authors of fake news. Most broadly, our results support the conclusion that encouraging people to engage in more thinking will be beneficial rather than harmful.

Acknowledgments

We gratefully acknowledge funding from the Ethics and Governance of Artificial Intelligence Initiative of the Miami Foundation (DGR and GP), the William and Flora Hewlett Foundation (DGR and GP), the Social Sciences and Humanities Research Council of Canada (GP), ANR grant ANR-17-EURE-0010 (Investissements d’Avenir program, BB), ANR Labex IAST (BB) and the Scientific Research Fund Flanders (FWO-Vlaanderen, BB).

Author note

Data, headlines and the supplementary materials are openly available at the project’s OSF page: osf.io/egy8p.

References

- Bago, B., & De Neys, W. (2017). Fast logic?: Examining the time course assumption of dual process theory. *Cognition*, *158*, 90–109.
- Bago, B., & De Neys, W. (2019). The intuitive greater good: Testing the corrective dual process model of moral cognition. *Journal of Experimental Psychology: General*, *148*(10), 1782–1801.
- Bronstein, M. V., Pennycook, G., Bear, A., Rand, D. G., & Cannon, T. D. (2019). Belief in fake news is associated with delusionality, dogmatism, religious fundamentalism, and reduced analytic thinking. *Journal of Applied Research in Memory and Cognition*, *8*(1), 108–117.
- Brown, L. D., Cai, T. T., & DasGupta, A. (2001). Interval estimation for a binomial proportion. *Statistical Science*, 101–117.
- Charness, G., & Dave, C. (2017). Confirmation bias with motivated beliefs. *Games and Economic Behavior*, *104*, 1–23.
- Dawson, E., Gilovich, T., & Regan, D. T. (2002). Motivated Reasoning and Performance on the was on Selection Task. *Personality and Social Psychology Bulletin*, *28*(10), 1379–1387.
- De Neys, W. (2006). Automatic–heuristic and executive–analytic processing during reasoning: Chronometric and dual-task considerations. *The Quarterly Journal of Experimental Psychology*, *59*(6), 1070–1100.
- De Neys, W., Novitskiy, N., Geeraerts, L., Ramautar, J., & Wagemans, J. (2011). Cognitive control and individual differences in economic ultimatum decision-making. *PloS One*, *6*(11), e27107.
- Ecker, U. K., & Ang, L. C. (2019). Political attitudes and the processing of misinformation corrections. *Political Psychology*, *40*(2), 241–260.

- Ecker, U. K., Lewandowsky, S., Swire, B., & Chang, D. (2011). Correcting false information in memory: Manipulating the strength of misinformation encoding and its retraction. *Psychonomic Bulletin & Review*, *18*(3), 570–578.
- Evans, J. (2010). *Thinking Twice: Two Minds in One Brain*. Oxford: Oxford University Press.
- Evans, J. S. B., & Stanovich, K. E. (2013). Dual-process theories of higher cognition advancing the debate. *Perspectives on Psychological Science*, *8*(3), 223–241.
- Frederick, S. (2005). Cognitive reflection and decision making. *The Journal of Economic Perspectives*, *19*(4), 25–42.
- Haidt, J. (2001). The emotional dog and its rational tail: a social intuitionist approach to moral judgment. *Psychological Review*, *108*(4), 814–834.
- Horton, J. J., Rand, D. G., & Zeckhauser, R. J. (2011). The online laboratory: Conducting experiments in a real labor market. *Experimental Economics*, *14*(3), 399–425.
- Johnson, H. M., & Seifert, C. M. (1994). Sources of the continued influence effect: When misinformation in memory affects later inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(6), 1420–1436.
- Kahan, D. (2018). Who “falls for” fake news? Apparently no one. Retrieved from <http://www.culturalcognition.net/blog/2018/10/25/who-falls-for-fake-news-apparently-no-one.html>
- Kahan, D. M. (2013). Ideology, motivated reasoning, and cognitive reflection. *Judgment and Decision Making*, *8*(4), 407–424.
- Kahan, D. M. (2017). *Misconceptions, misinformation, and the logic of identity-protective cognition*. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2973067
- Kahan, D. M., Peters, E., Dawson, E. C., & Slovic, P. (2017). Motivated numeracy and enlightened self-government. *Behavioural Public Policy*, *1*(1), 54–86.

- Kahan, D. M., Peters, E., Wittlin, M., Slovic, P., Ouellette, L. L., Braman, D., & Mandel, G. (2012). The polarizing impact of science literacy and numeracy on perceived climate change risks. *Nature Climate Change*, 2(10), 732–735.
- Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Straus and Giroux.
- Lazer, D. M., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., ... Rothschild, D. (2018). The science of fake news. *Science*, 359(6380), 1094–1096.
- Lewandowsky, S., Ecker, U. K., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest*, 13(3), 106–131.
- Lewandowsky, S., Gignac, G. E., & Vaughan, S. (2013). The pivotal role of perceived scientific consensus in acceptance of science. *Nature Climate Change*, 3(4), 399–404.
- Martel, C., Pennycook, G., & Rand, D. (2019). *Reliance on emotion promotes belief in fake news*. Retrieved from <https://psyarxiv.com/a2ydw/>
- Mercier, H., & Sperber, D. (2011). Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences*, 34(02), 57–74.
- Miyake, A., Friedman, N. P., Rettinger, D. A., Shah, P., & Hegarty, M. (2001). How are visuospatial working memory, executive functioning, and spatial abilities related? A latent-variable analysis. *Journal of Experimental Psychology: General*, 130(4), 621–640.
- 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*, 188, 39–50.
- Pennycook, G., & Rand, D. G. (2019b). Who falls for fake news? The roles of bullshit receptivity, overclaiming, familiarity, and analytic thinking. *Journal of Personality*. <https://doi.org/10.1111/jopy.12476>

- Roozenbeek, J., & van der Linden, S. (2019a). Fake news game confers psychological resistance against online misinformation. *Palgrave Communications*, 5(1), 12.
- Roozenbeek, J., & van der Linden, S. (2019b). The fake news game: actively inoculating against the risk of misinformation. *Journal of Risk Research*, 22(5), 570–580.
- Shtulman, A., & McCallum, K. (2014). Cognitive reflection predicts science understanding. *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*, 2937–2942.
- Signorell, A. (2016). DescTools: Tools for descriptive statistics. *R Package Version 0.99*, 17.
- Sloman, S. A., & Rabb, N. (2019). Thought as a determinant of political opinion. *Cognition*, 188, 1–7.
- Stanovich, K. (2011). *Rationality and the reflective mind*. Oxford: Oxford University Press.
- Stewart, N., Chater, N., & Brown, G. D. (2006). Decision by sampling. *Cognitive Psychology*, 53(1), 1–26.
- Swami, V., Voracek, M., Stieger, S., Tran, U. S., & Furnham, A. (2014). Analytic thinking reduces belief in conspiracy theories. *Cognition*, 133(3), 572–585.
- Swire, B., Berinsky, A. J., Lewandowsky, S., & Ecker, U. K. (2017). Processing political misinformation: comprehending the Trump phenomenon. *Royal Society Open Science*, 4(3), 160802.
- Thompson, V. A., Turner, J. A. P., & Pennycook, G. (2011). Intuition, reason, and metacognition. *Cognitive Psychology*, 63(3), 107–140.
- Thomson, K. S., & Oppenheimer, D. M. (2016). Investigating an alternate form of the cognitive reflection test. *Judgment and Decision Making*, 11(1), 99–113.
- Trémolière, B., De Neys, W., & Bonnefon, J.-F. (2012). Mortality salience and morality: Thinking about death makes people less utilitarian. *Cognition*, 124(3), 379–384.
- van der Linden, S., Leiserowitz, A., & Maibach, E. (2018). Scientific agreement can neutralize politicization of facts. *Nature Human Behaviour*, 2(1), 2–3.

