

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Evaluating Information and Misinformation during the COVID-19 Pandemic: Evidence for Epistemic Vigilance

Permalink

<https://escholarship.org/uc/item/4nq6t2d3>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

Authors

Vasilyeva, Nadya
Smith, Kristopher M
Barr, Kelli
[et al.](#)

Publication Date

2021

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Evaluating Information and Misinformation during the COVID-19 Pandemic: Evidence for Epistemic Vigilance

Nadya Vasilyeva^{1,2}; Kristopher M. Smith³, Kelli Barr⁴, Jordan Kiper⁵, Stephen Stich⁶, Edouard Machery⁴, H. Clark Barrett¹

¹Department of Anthropology, University of California, Los Angeles

²Department of Psychology, California State University East Bay

³Department of Psychology, University of Pennsylvania

⁴Department of History and Philosophy of Science, University of Pittsburgh

⁵Department of Anthropology, University of Alabama

⁶Philosophy Department, Rutgers University

Abstract

There are many ways to go wrong when evaluating new information, e.g. by putting unwarranted trust in non-experts, or failing to scrutinize information about threats. We examined how effective people were at evaluating information about the COVID-19 pandemic. Early in the course of the pandemic, we recruited 1791 participants from six countries with varying levels of pandemic severity, and asked them to evaluate true and false pandemic-related statements (assertions and prescriptions) sampled from the media. We experimentally manipulated the source of each statement (a doctor, a political/religious leader, social media, etc.). Overall, people proved to be epistemically vigilant: they distinguished between true and false statements, especially prescriptions, and they trusted doctors more than other sources. These effects were moderated by feeling threatened by the pandemic, and by strong identification with some sources (political/religious leaders). These findings provide optimism in the fight against misinformation, while highlighting challenges posed by politics and ideology.

Keywords: epistemic vigilance; trust; misinformation; COVID-19

We are at a moment of historically unprecedented concern about how true and false information spreads through communities and around the world. This concern took center stage with the arrival of the global COVID-19 pandemic in early 2020. This pandemic has brought the dangers of false belief contagion to the fore, increasing concerns about the roles of social media and politics in disrupting public health messaging. In a very real sense, COVID-19 has been both a medical pandemic and a misinformation pandemic as well (Tagliabue et al., 2020).

Soon after the beginning of the pandemic, in May 2020, we conducted a study examining how various social and personal factors impact peoples' ability to evaluate true and false information about the pandemic. The majority of what laypeople know about the pandemic—for example, about the efficacy of masks, mechanisms of disease transmission, and procedures to guard against infection—is learned from other people (as opposed to personal experience). Social learning, or the acquisition of information from others, is notoriously prone to error, because it involves trust (Boyd & Richerson, 1985). Many forms of social learning rely on the assumption that information comes from knowledgeable, helpful and

cooperative teachers who selectively convey information that is true and maximally useful to the learners. Needless to say, the reality does not always live up to this expectation. Successful learners thus need to engage in *epistemic vigilance*: the cognitive processes of evaluation and scrutiny that individuals and communities use to determine whether to trust or distrust a given testimony, based on its source and content (Landrum et al., 2015; Sperber et al., 2010). With proper settings, epistemic vigilance should guard recipients against believing false claims, helping them hone in on true information, although it does not guarantee it. The specific information evaluation criteria, as well as the overall propensity to (dis)trust, are likely to play a role in how information and misinformation spread (O'Connor & Weatherall, 2019). A universally relevant, widely discussed and contested topic, such as an ongoing pandemic, presents a particularly telling opportunity to examine epistemic vigilance in action. Under the conditions of having to sort through a sea of potentially relevant but often conflicting information, selective endorsement of accurate beliefs can be taken as evidence of “well-tuned” epistemic vigilance processes, enabling a balance between being overly trustful and being overly skeptical.

Theorists of cultural transmission have identified a variety of social factors that are likely to influence the uptake and spread of information, such as the source of information and the relationship between the recipient and source (including individuals, groups, and media) and information content (Landrum et al., 2015; Sperber et al., 2010). For example, people might deem doctors generally trustworthy sources of medical information, but could discount a doctor's claims if contradicted by a trusted political leader. The source- and content-based assessments are likely to interact in a variety of ways, from mutually reinforcing to mutually cancelling, complicating the link between epistemic vigilance and accuracy of beliefs. E.g., people might exercise epistemic vigilance by checking novel claims for internal coherence and/or consistency with prior beliefs, but suspend, override – or, on the contrary, intensify – such scrutiny depending on the source of the claim. Implausible claims might become more plausible if heard from a trusted source, or if supported by many others in one's social network.

Information content can also have direct effects. For example, negative information is more likely to be attended to, believed, and remembered than positive information (Baumeister et al., 2001), possibly because the costs of disbelieving information about potential threats is greater than the benefit of ignoring false information about threats (Fessler et al., 2014). The asymmetry in costs and benefits of hazardous information is thought to bias people toward epistemic trust and greater credulity toward this information. Subjective perception of threat might further make information about hazards in the environment more relevant, such that people who perceive themselves particularly at risk may lower their threshold to entertain something as true, risking a false positive in order to avoid a costly miss (Fessler et al., 2017; Haselton & Nettle, 2006).

Research on the spread of misinformation points to reasons to be pessimistic about the capacity of epistemic vigilance to ward off false beliefs. While prior work suggests that people do often attend to cues of consensus—including scientific consensus on matters such as climate change (Lewandowsky et al., 2012)—a variety of factors can disrupt this, including motivated cognition (rejecting findings that threaten core beliefs or worldviews) and “echo chambers” that result from self-assortment of information-sharing communities, especially online (Kahan et al., 2011; Linden et al., 2017; Vicario et al., 2016). Peoples’ perceptions of risks are shaped by worldviews, political orientations, and other cultural factors (D.M. Fessler et al., 2017; Kahan et al., 2011). Cues from political leaders can sway peoples’ opinions away from scientific consensus (Brulle et al., 2012). Sadly, there is evidence that education and scientific literacy can actually increase rather than decrease partisanship in views of science (Lewandowsky & Oberauer, 2016). While none of this may be surprising in 2021, it suggests that we might expect peoples’ evaluations of statements related to the pandemic to be influenced by the source of the information, their political orientations, attitudes towards science, and perceived personal risk from the pandemic. And, given the contested nature of some information about the pandemic—for example, widespread claims by prominent political leaders about the efficacy of false cures, and skepticism about masks—we might expect to see confusion in peoples’ judgments about what is true and false in this domain.

We examined how these and other factors shaped people’s response to information about the coronavirus pandemic. One objective was to examine how good people were at discriminating true from false information about coronavirus, during the pandemic. Another important goal was to examine how the *source* of information—for example, doctors, politicians, or social media—influences peoples’ ability to discriminate true from false information. The third goal was to examine how the source effects varied across peoples’ personal attributes, such as their ideological orientations, and the extent to which they felt threatened by the pandemic. Finally, we sought to examine possible differences across six countries that varied, at the time of the study, in how severely they were being hit by the pandemic: the U.S., Spain, and

Ecuador (relatively high pandemic severity in May 2020), and New Zealand, Germany, and Sweden (relatively low pandemic severity at the time).

We presented participants with a set of real-world, fact-checked statements and experimentally manipulated the source of the statements. Each statement was attributed to a doctor, a religious leader, a political leader from one’s own or the opposition party, or presented as posted on social media; we also had a set of statements from unspecified sources as controls. Additionally, we measured person-level descriptors, including demographic characteristics, beliefs, and ideological orientations. We were thus able to examine whether some sources, such as doctors and political leaders, are deemed universally trustworthy, whether the capacity to discriminate true and false claims varied across sources, and whether these effects were moderated by individual-level attributes.

In the face of a real-world, global threat to their health and safety, how do people balance epistemic trust and vigilance in an environment rife with conflicting information? If epistemic vigilance is undermined by concerns of group loyalty and ideology, then we would expect people to trust sources that align with their group membership more than a domain expert. And if threat biases people toward epistemic trust, we would expect people who perceive themselves more at risk to be more credulous toward pandemic-related information, regardless of whether the information is true. On the other hand, if epistemic vigilance is effective, we would expect people to be effective at distinguishing true or false propositions. We would also expect people to trust the doctor as a domain expert more than respected figures in irrelevant domains such as religion and politics.

Method

Participants

Participants ($N=1791$; age $M = 44$ years, $SD = 15$, range 18–88) from USA, Ecuador, Spain, Sweden, Germany and New Zealand were recruited through Qualtrics panel. Data collection took place over 20 days in a relatively early stage of the COVID-19 pandemic (May 1 – 21, 2020). Additional participant information is available in the Online Supplement.

Materials, Design and Procedure

Each participant was presented with 36 medically-relevant statements and rated their confidence that the statement is true or false. All statements were real claims that appeared in the international media. The statements were fact-checked (see OS for more details) and tagged as either true (supported by the evidence) or false (unsupported by the evidence). Half of the statements were true, and half were false. We will call this factor *truth*. Twelve of the statements were coronavirus related assertions (“COVID facts”), twelve were coronavirus-related recommendations (“COVID prescriptions”), and 12 were medical assertions unrelated to coronavirus (“non-COVID facts”)—we call this *statement type*. Some statements were slightly edited to ensure clarity,

similar average length across statement types, and/or fact- vs. prescription-format. Sample statements are shown in Table 1. The complete list of statements and additional details on constructing the stimulus set are available in the OS.

Each participant saw six blocks containing six statements each. All statements in a given block were attributed to one of six “sources”: a *medical doctor*, a *self-aligned political leader* (a political leader representing the party or coalition that the participant identified with most), an *opposition political leader* (a political leader representing the party that the participant specified as the most obvious and prominent “opposition party” to their own), a *religious leader* of the participant’s faith (or, if non-religious, a religious leader of their family faith or the mainstream religion in their region), or a statement posted on *social media*, or a statement from an *unspecified source* (presented in isolation, not attributed to any source). Specific content for political and religious leaders was determined by a participant’s earlier responses about their political affiliation and religiosity.

For example, on a trial attributing a false COVID prescription to a doctor, participants might read: “A medical doctor made the following statement: [statement text]. Rate your confidence about this statement: [statement text]”; this was followed by a rating scale ranging from 1 “very confident it is false” to 5 “very confident it is true”, with the midpoint labeled as “not confident whether it is true or false”. We will refer to higher levels of confidence that a statement is true as greater *endorsement* of the statement.

Specific statements rotated through sources following a counterbalancing scheme which ensured that across participants, all statement-source combinations occurred with similar frequency. The order of source blocks was randomized for each participant. In sum, the study had a 2

(truth: true or false) $\times$ 3 (statement type: COVID prescription, COVID fact, or non-COVID medical fact) $\times$ 6 (source: doctor, self-aligned political leader, opposition political leader, religious leader, social media, or unspecified) within-subjects design.

To explore whether the effects of experimental manipulations interacted with participants’ characteristics, we included several additional measures. First, participants evaluated the subjective level of threat from the coronavirus pandemic, by indicating how threatened they personally felt, on a scale from 1 (not at all threatened) to 5 (extremely threatened). Second, participants rated their agreement with five statements measuring their endorsement of science, e.g. “The scientific method we have is the best method we have for learning the truth about how the world works” (1 strongly disagree – 5 strongly agree; one question reverse-scored); responses were combined into a single weighted factor score “scientism”. Third, participants rated the importance of religion in their life (1 not at all important – 5 completely important). Finally, participants rated their political orientation (1 very much on the left / very liberal – 5 very much on the right / very conservative), and indicated the strength of affiliation with political party or coalition with which they identify most (1 I barely identify with this party at all – 5 Extremely strongly).

Additional demographic and supplementary measures are described in the OS. The study was pre-registered at aspredicted.org/be9g9.pdf; some analyses presented here were exploratory and are clearly labeled as such below. All materials and surveys are available at https://osf.io/xhsr/?view_only=aa0ef8a218fa47628b1d86091e6b25f8.

Results

Analytic approach Statement endorsement ratings were treated as an ordinal variable, and analyzed in a series of regression models implemented using *clmm* command from the *ordinal* R package, with flexible thresholds and random intercepts for participants, items (individual statements), and country (unless stated otherwise). We present model comparison statistics between models with and without the effect of interest and the relevant coefficient statistics for models that significantly improved fit. Specific comparisons for predictors with more than two levels were performed by releveling and querying the same model.

Effects of statement truth and type Truth of a statement strongly predicted its endorsement, $\chi^2(1) = 23.12, p < 0.001$: true statements were endorsed more than false statements ($M_{\text{true}} = 3.64, M_{\text{false}} = 2.44$). This effect varied across statement types, $\chi^2(2) = 10.46, p = 0.005$, although statement type did not independently predict endorsement, $\chi^2(2) = 4.30, p = 0.116$. As Figure 1 shows, people tended to discriminate between true and false statements more reliably for COVID prescriptions than for COVID facts, $b = 1.31, z = 1.957, p = 0.050$, and Non-COVID facts $b = 2.40, z = 3.54, p < 0.001$, although the former difference was marginal; COVID facts showed a trend for a more reliable true vs. false

Table 1: Sample statements.

True	False
<u>COVID fact</u>	
A child infected with coronavirus is less likely to show symptoms (fever, cough, shortness of breath), or will develop milder symptoms, than an infected adult.	Vaccines against pneumonia can protect you against COVID-19 by triggering production of the antibodies that attack coronavirus.
<u>COVID prescription</u>	
You should limit the time children spend with elderly adults and people with serious medical conditions who are at high risk for severe illness from COVID-19.	You should avoid mosquitoes because they can transmit coronavirus from person to person.
<u>Non-COVID fact</u>	
Some heart attacks happen without any pain whatsoever.	People who drink milk or other dairy products when they are sick produce more nasal mucus.

discrimination than Non-COVID facts, but the difference did not reach significance, $b = 1.07, z = 1.59, p = 0.112$.

Effect of statement source The source to which statements were attributed also significantly affected their endorsement, $\chi^2(5) = 380.68, p < 0.001$. Statements attributed to doctors were endorsed the most ($M = 3.23$, all pairwise $ps < 0.001$), followed by statements from unspecified sources and self-aligned political leaders, which did not differ from each other ($Ms = 3.05$ and 3.04 respectively, $p = 0.463$). These in turn were endorsed more ($ps < 0.001$) than statements from opposition political leaders, religious leaders, and social media, which did not differ from each other ($M = 2.98, 2.98, 2.97$, respectively; $ps \geq 0.470$). The source effect did not interact with statement truth or type, $\chi^2(10) = 13.88, p = 0.178$. Since specific statements rotated through all sources across subjects, the effect of statement source is not attributable to statement content. Of these effects, the privileged status of a doctor as a source was predicted, as well as the difference between self-aligned and opposition political leader. Contrary to predictions, social media did not emerge as the least trusted source. Statements from unspecified sources were endorsed highly, yielding only to statements from doctors and aligned political leaders, suggesting a fairly high baseline level of trust in information not marked by source.

Source effects were not moderated by statement truth, $\chi^2(5) = 7.25, p = 0.201$, or statement content, $\chi^2(10) = 13.25, p = 0.210$. That is, regardless of whether the statement was true or false, participants were more likely to believe a doctor than any other source. Finally, the three-way interaction between statement truth, source and type was not significant, $\chi^2(10) = 9.74, p = 0.463$.




Figure 1. Distribution of confidence ratings by statement truth and type. The x-axis shows the response categories to the confidence question and the y-axis plots the probability of that response.

Interactions with individual-level predictors

Demographics We analyzed whether endorsement of statements varied as a function of demographic variables, specifically gender, age, education, and economic status (all standardized), and the interaction between these variables and truth, running separate models for each demographic variable. Gender did not significantly relate to endorsement, $b = 0.03, z = 0.87, p = 0.384$, and did not significantly interact with truth, $\chi^2(1) = 1.95, p = 0.163$. Age significantly interacted with truth, $\chi^2(1) = 137.32, p < 0.001$; older participants were less likely to mistakenly rate false statements as true, $b = -0.20, z = -12.14, p < 0.001$, but endorsement of true statements did not vary with age, $b = -0.03, z = -1.79, p = 0.074$. Greater education was linked to higher endorsement of statements in general, $b = 0.04, z = 2.38, p = 0.017$, equally for true and false claims, $\chi^2(1) = 0.03, p = 0.856$. Participants reporting higher economic status were also more likely to classify statements as true, and this effect was stronger for true ($b = 0.07, z = 4.42, p < 0.001$) than false statements ($b = 0.04, z = 2.29, p = 0.022$; interaction $\chi^2(1) = 5.79, p = 0.016$). To assess unique contributions of each demographic variable, multiple variables were entered as predictors at once. Adding economic status (with status x truth interaction) significantly improved the fit over the model with age (and age x truth interaction) only, $\chi^2(2) = 20.71, p < 0.001$. Likewise, including age improved the model over economic status only, $\chi^2(2) = 199.10, p < 0.001$; however, including education into the model with age and economic status did not improve fit, $\chi^2(1) = 1.56, p = 0.212$.

Thus, the best fitting model of demographics included the age x truth and economic status x truth interactions; the subsequent analyses on other individual predictors control for these effects.

Perceived threat There was a three-way interaction between participants' perceived threat from COVID, statement type and truth ($\chi^2(2) = 28.13, p < 0.001$; see Figure 2). In general, participants who felt more threatened were more likely to rate statements as true. This effect was present for true and false statements of each type; however, it was more pronounced for true COVID prescriptions, $b = 0.20, z = 9.45, p < 0.001$, than for all the other statement types, true or false (false non-COVID fact, $b = 0.06, z = 2.96, p = 0.003$; false COVID fact, $b = 0.05, z = 2.49, p = 0.013$; false COVID prescription, $b = 0.05, z = 2.66, p = 0.008$; true non-COVID fact, $b = 0.05, z = 2.37, p = 0.018$; true COVID fact, $b = 0.06, z = 2.84, p = 0.005$). There was no significant four-way interaction with statement source, $p = 0.635$.

Scientism Participants' with higher levels of endorsement of science ('scientism') were less likely to endorse false statements as true, $b = -0.13, z = 7.34, p < 0.001$, and were more likely to endorse true statements as true, $b = 0.20, z = 10.65, p < 0.001$ (interaction $\chi^2(1) = 386.34, p < 0.001$). Greater scientism was also selectively associated with stronger endorsement of statements coming from a doctor, $b = 0.08, z = 3.10, p = 0.002$ (scientism x source $\chi^2(5) = 13.79$,




Figure 2. Distribution of confidence ratings by statement truth, type, and participant's perceived threat from coronavirus. The x-axis shows the response categories to the confidence question and the y-axis plots the probability of that response.

$p = 0.017$). The three-way interaction between truth, source, and scientism was nonsignificant, $\chi^2(10) = 11.89, p = 0.292$.

Religious importance Participants who placed greater importance on religion were more likely to endorse false statements as true, $b = 0.15, z = 12.58, p < 0.001$. The effect of religious importance on endorsement of true statements was significantly less pronounced, $b = 0.03, z = 2.78, p = 0.005$ (interaction $\chi^2(1) = 127.17, p < 0.001$). Religious importance also moderated source effects, $\chi^2(5) = 25.42, p = 0.001$. Participants who placed greater importance on religion were more likely to endorse statements from a religious leader as true, compared to less religious participants, $b = 0.14, z = 8.74, p < 0.001$. Endorsement of claims from other sources was not as affected by participant religiosity, $ps < 0.05$. The three-way interaction between truth, source, and religious importance was nonsignificant, $\chi^2(10) = 11.39, p = 0.328$.

Strength of Political identification

The extent to which participants identified with their political party was related to their discrimination between true and false claims, $\chi^2(1) = 19.88, p < 0.001$. Participants who identified more with their party were more likely to endorse true statements, $b = 0.07, z = 4.10, p < 0.001$, but were not more likely to endorse false statements, $b = 0.00, z = 0.05, p = 0.957$. Strong party identification was also associated with endorsing statements coming from a self-aligned political leader as true, $b = 0.09, z = 3.86, p < 0.001$, but did not predict endorsement of statements coming from other sources, $ps >$

0.05 (interaction, $\chi^2(5) = 16.23, p = 0.006$). There was no three-way interaction, $\chi^2(10) = 15.58, p = 0.112$.

Political orientation High political conservatism was associated with increased erroneous endorsement of false statements, $b = 0.09, z = 5.45, p < 0.001$; in contrast, endorsement of true statements was unrelated to political conservatism, $b = 0.03, z = 1.92, p = 0.055$ (interaction $\chi^2(1) = 14.27, p < 0.001$). Political orientation did not interact with source, $\chi^2(5) = 6.68, p = 0.246$, and the three-way interaction was not significant, $\chi^2(15) = 16.56, p = 0.346$.

Country effects Finally, we examined pandemic severity as a country-level variable (this analysis omitted random country intercepts). Our *a priori* classification of countries in terms of high vs. low severity was confirmed by mean ratings of perceived threat from COVID ($M_{hi}=3.25, M_{lo}=2.57; b=1.16, \chi^2(1)=16.56, p < 0.001$, treating threat ratings as an ordinal outcome). People in high severity countries were worse at discriminating between true and false statements, $\chi^2(1)=48.31, p < 0.001$, due to higher endorsement of false statements, resulting in higher average statement endorsement in high severity countries, $\chi^2(1)=51.00, p < 0.001$. This effect was driven by the participants in the US and Ecuador who rated false statements about COVID higher than the other countries ($ps < 0.001$). On average, statement endorsement ratings were the highest in the U.S. (all $ps < 0.001$), followed by Ecuador (all $ps \leq 0.002$), followed by Germany, whose ratings were in turn higher than Sweden and Spain ($ps .013$ and $.010$). New Zealand did not significantly differ from

Germany, Sweden or Spain ($p=.128$, $.350$ and $.291$), and the latter two did not differ from each other ($p=.900$).

Discussion and conclusion

Our findings are particularly interesting in light of prior literature on how information and misinformation spread, and in light of the alarm over the misinformation pandemic related to COVID-19 in particular (Tagliabue et al., 2020). Particularly noteworthy is our main finding that people were good at distinguishing true from false information in this domain—which might be surprising given the highly politicized and contested nature of information about the pandemic. In fact, people were best at distinguishing true from false prescriptions about how to behave during the pandemic, suggesting that they were indeed tracking medical consensus at the time, and were particularly focused on actionable information about what to do.

Combined with the finding that doctors were the most trusted source overall, we consider this good news. Epistemic vigilance appears to have been operating well during this emergency, at least for this set of participants and materials. We suspect that participants' accuracy was bolstered by the fact that discussions of information about the virus were in heavy circulation at the time, possibly increasing peoples' degree of scrutiny and epistemic vigilance.

That said, we did find evidence of small biasing effects acting on top of overall accurate epistemic evaluations. Some of them were predicted and are consistent with other literature. For example, after doctors, the second-most trusted source overall (on par with unsourced information) was self-aligned political leaders, consistent with prior findings that politicians can sway people away from attending to scientific consensus (Brulle et al., 2012). This pattern is particularly concerning in countries where politicians from different fractions disagree on matters of science, such as deadliness of a virus or reality of climate change. This generally shows a downside of epistemic trust, when people trust "experts" who are not necessarily appropriate to a particular domain. We also found that personal orientation towards science enhanced peoples' epistemic trust of doctors, which is reassuring, though it indicates that evaluation of scientific information is shaped by upbringing and ideology.

We also found, as expected, that perceived personal threat from the pandemic influenced peoples' epistemic vigilance, but in a different way than we expected. Peoples' self-rated risk actually increased relative endorsement of true prescriptions. This was in contrast to our original pessimistic prediction that threat would in general give a stronger credibility boost to false rather than true statements. It is not straightforward how our pattern of results could be reconciled with theories that argue asymmetries in the costs and benefits of credulity toward hazardous information bias people toward believing all threat-relevant information (Fessler et al., 2014; 2017). The finding of threat-related improvement in discriminating true from false guidelines for actions that

mitigate one's risk corroborates recent arguments that people are better able to distinguish between information and misinformation when motivated to do so (Pennycook & Rand, 2021); people who consider themselves at high risk of COVID may be more motivated to identify useful information to prevent contracting or spreading the pathogen.

Our study was conducted across six countries, and the key reported effects hold across these countries. However, we did observe some country differences that are noteworthy. In particular, the U.S. and Ecuador showed the highest degree of epistemic trust overall—with the downside that people in these countries were more credulous of false COVID facts and prescriptions than elsewhere. These were among the three countries we selected because of the heavy impact COVID had already taken at the time (along with Spain, which did not show the same credulity patterns). This country-level differences in credulity are intriguing, but attributing it to specific cultural and political differences or the local pandemic situation is beyond the scope of this paper.

While we examined generalizability across several countries, we only examined information from a medical domain and we do not know if our design would have produced similar results for a different threat, such as climate change. Our results suggest that people are less epistemically vigilant toward less immediate threats that are perceived to pose less personal risk. Thus, threats such as climate change may be more susceptible to misinformation. Our selection of specific stimuli may further limit the generalizability of the results; the specific statements may have been unusually salient, making it easier for participants to distinguish between true and false statements. The effect of epistemically suspect sources or factors may in fact be larger with a different set of stimuli.

A virtue of our study is that we used real-world statements from the media environment fact-checked by trusted authorities. This allowed us to estimate, at least for the tested materials, how effectively people distinguish between true and false information about the pandemic. This also provided some ecological validity for our study: these are actual statements that participants have to evaluate to potentially guide their behavior.

In summary, we found that people were good at discriminating true from false information about the pandemic, during the pandemic. This is good news, and suggests that the plausibly present forces of misinformation and disinformation were not enough to derail epistemic evaluation in this domain. Moreover, subjective threat from the pandemic was associated with an improved ability to discriminate true from false information in this domain. And when people did use the source of a proposition to evaluate whether a statement was true, they did so—mostly—in an epistemically justifiable way, trusting a domain expert more than other source. However, we did find evidence for the kinds of processes thought to work against proper epistemic evaluation, including effects of politics and group identity on orientation towards sources, as well as effects of attitudes towards or against science and religion.

Taken together, these findings provide grounds for optimism in the fight against misinformation, but they also reinforce the challenges posed by politics, ideology, education, and degrees of risk people perceive in disasters and emergencies.

References

- Baumeister, R.F., Bratslavsky, E., Finkenaur, C., & Vohs, K.D. (2001). Bad is stronger than good. *Review of General Psychology*, 5(4), 323-370.
- Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago Press.
- Brulle, R. J., Carmichael, J., & Jenkins, J. C. (2012). Shifting public opinion on climate change: An empirical assessment of factors influencing concern over climate change in the U.S., 2002–2010. *Climatic Change*, 114(2), 169–188. <https://doi.org/10.1007/s10584-012-0403-y>
- Fessler, D.M., Pisor, A. C., & Holbrook, C. (2017). Political orientation predicts credulity regarding putative hazards. *Psychological Science*, 28(5), 651–660.
- Fessler, D.M., Pisor, A. C., & Navarrete, C. D. (2014). Negatively-biased credulity and the cultural evolution of beliefs. *PloS One*, 9, 95167.
- Haselton, M. G., & Nettle, D. (2006). The paranoid optimist: An integrative evolutionary model of cognitive biases. *Personality and Social Psychology Review* 10(1), 47–66. https://doi.org/10.1207/s15327957pspr1001_3
- Kahan, D. M., Jenkins-Smith, H., & Braman, D. (2011). Cultural cognition of scientific consensus. *Journal of Risk Research*, 14(2), 147–174. <https://doi.org/10.1080/13669877.2010.511246>
- Landrum, A. R., Eaves Jr, B. S., & Shafto, P. (2015). Learning to trust and trusting to learn: A theoretical framework. *Trends in Cognitive Sciences*, 19(3), 109–111.
- Lewandowsky, S., Gignac, G.E., & Vaughn, S. (2012). The pivotal role of perceived scientific consensus in acceptance of science. *Nature Climate Change* 3: 399–404
- Lewandowsky, S., & Oberauer, K. (2016). Motivated Rejection of Science. *Current Directions in Psychological Science*, 25(4), 217–222. <https://doi.org/10.1177/0963721416654436>
- Linden, S. van der, Leiserowitz, A., Rosenthal, S., & Maibach, E. (2017). Inoculating the Public against Misinformation about Climate Change. *Global Challenges*, 1(2), 1600008. <https://doi.org/10.1002/gch2.201600008>
- O'Connor, C., & Weatherall, J. O. (2019). *The misinformation age: How false beliefs spread*. Yale University Press.
- Pennycook, G., & Rand, D.G., (2021). The psychology of fake news. *Trends in Cognitive Sciences*, 25(5), 388–402.
- Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origi, G., & Wilson, D. (2010). Epistemic vigilance. *Mind & Language*, 25(4), 359–393.
- Tagliabue, F., Galassi, L., & Mariani, P. (2020). The “Pandemic” of Disinformation in COVID-19. *Sn Comprehensive Clinical Medicine*, 1–3. <https://doi.org/10.1007/s42399-020-00439-1>
- Vicario, M. D., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., Stanley, H. E., & Quattrociocchi, W. (2016). The spreading of misinformation online. *Proceedings of the National Academy of Sciences*, 113(3), 554–559. <https://doi.org/10.1073/pnas.1517441113>