



ARTICLE



<https://doi.org/10.1057/s41599-023-02329-y>

OPEN

How do online users respond to crowdsourced fact-checking?

Folco Panizza¹ , Piero Ronzani², Tiffany Morisseau³, Simone Mattavelli⁴ & Carlo Martini^{5,6}

Recently, crowdsourcing has been proposed as a tool for fighting misinformation online. Will internet users listen to crowdsourced fact-checking, and how? In this experiment we test how participants follow others' opinions to evaluate the validity of a science-themed Facebook post and examine which factors mediate the use of this information. Participants observed a post presenting either scientific information or misinformation, along with a graphical summary of previous participants' judgements. Even though most participants reported not having used information from previous raters, their responses were influenced by previous assessments. This happened regardless of whether prior judgements were accurate or misleading. Presenting crowdsourced fact-checking however did not translate into the blind copying of the majority response. Rather, participants tended to use this social information as a cue to guide their response, while also relying on individual evaluation and research for extra information. These results highlight the role of individual reasoning when evaluating online information, while pointing to the potential benefit of crowd-sourcing-based solutions in making online users more resilient to misinformation.

¹ Molecular Mind Laboratory, IMT School for Advanced Studies Lucca, Lucca, Italy. ² International Security and Development Center, Berlin, Germany. ³ Université Paris Cité and Univ. Gustave Eiffel, LaPEA, Boulogne-Billancourt, France. ⁴ Department of Psychology, Bicocca University, Milano, Italy. ⁵ Centre for Applied and Experimental Epistemology, Department of Philosophy, Vita-Salute San Raffaele University, Milan, Italy. ⁶ Centre for Philosophy of Social Science, Department of Political and Economic Studies, University of Helsinki, Helsinki, Finland. ✉email: folco.panizza@imtlucca.it

Introduction

Social media's information landscape is still a mystery in many respects. While we have substantive knowledge of the type and quality of news disseminated by media outlets (Acerbi et al. 2022), little is known about the content generated by online users themselves, with the risk of false or misleading information being spread in an uncontrolled manner (Cook et al. 2015). Indeed, misinformation often goes unchecked in private communications through messaging apps or closed groups (Rossini et al. 2021; Tandoc, 2021). The consequences of this misguided communication are deeply worrying: bogus cures spread unhampered, damaging the health of those who take them (Álvarez-Moreno et al. 2021; Branch et al. 2022); climate action is curbed by misleading claims about the effectiveness of policies, if not the very existence of anthropogenic global warming (Coan et al. 2021; Lewandowsky, 2021).

In those contexts where professional fact-checkers cannot readily intervene in the conversation, several user-centred approaches have been tested to improve the quality of information (Fazio, 2020; Lewandowsky and Van Der Linden, 2021; Lorenz-Spreen et al. 2021; Panizza et al. 2022; Pennycook et al. 2021; Roozenbeek et al. 2023; Tully et al. 2021). Those approaches rely on nudging – namely, changing the context of the decision to “move people in directions that will make their lives better” (Thaler and Sunstein, 2009, p. 6), and on boosting – namely, targeting the competences of users (Hertwig and Grüne-Yanoff, 2017). Both approaches have devoted significant attention to interventions that provide information to stimulate users' critical thinking, like informative nudges (Acquisti et al. 2017) or educational nudges (Sunstein, 2016). One such intervention is to crowdsource fact-checking. Crowd-sourced fact-checking consists in recruiting internet users to evaluate information circulating online (Wojcik et al. 2022). In principle, crowdsourcing provides a rapid and far-reaching response to chain letters and viral content. Peer-to-peer communication can indeed be quite effective compared to communications from official sources (Druckman et al. 2018), and several studies show that when users are part of well-sorted groups, they are able to identify the quality of news outlets and their output (Martel et al. 2023). Social media companies have considered the potential of this approach: in 2021, Twitter introduced an experimental programme called Birdwatch (Coleman, 2021; Wojcik et al. 2022); likewise, Facebook CEO Mark Zuckerberg discussed the possibility of introducing crowdsourcing on Facebook (Zuckerberg and Zittrain, 2019).

One little-debated issue related to crowdsourced fact-checking is how lay ratings are received by peers. Namely, how do users use this additional information? What is the relative weight of such information in the decision process? Understanding the actual impact of crowd-sourced fact-checking could inform both social media platforms on whether or not to adopt it, and could guide policy to incentivise and regulate the process. Evidence suggests that many online behaviours are socially influenced: people's opinions are swayed by previous online recommendations (Zhu et al. 2012); financial decisions can be affected by the desire for group cohesion in online communities (Breitsohl et al. 2015); different opinions tend to converge on Wikipedia's discussion pages, regardless of the degree of anonymity of the participating members (Tsikerdekis, 2013), suggesting that even complete strangers could influence each other. Evaluating content on one's own news feed might also be informed by others' opinions: negative comments to a Facebook post can reduce the perceived persuasiveness of linked articles (Winter et al. 2015) and even simple refutations by other users can improve the recognition of false content (Colliander, 2019).

While this evidence suggests that providing crowdsourced fact-checking could significantly affect users' evaluations of news

content, few studies have directly examined users' response to this approach to the best of our knowledge. According to the MAIN model on the credibility of sources (Sundar, 2008), an environmental change in the decision setting, such as the presentation of others' fact-checking might trigger heuristic reasoning in users and ultimately influence their beliefs and decisions. One heuristic that may be generated by the evaluation of other users might be that “in the absence of other cues, others' opinions should be trusted.” Yet, this heuristic is not obvious, especially considering that several studies unveiled that individuals tend to believe that others, rather than themselves, are the most susceptible to the spread of misinformation (Altay and Acerbi, 2023; Jang and Kim, 2018; Ștefăniță et al., 2018; Yoo et al. 2022). In other words, how can others' opinions be trusted if people believe that others are generally more gullible than themselves?

A series of surveys and a field experiment on Twitter suggest that recruiting users to add context to a potentially misleading tweet can reduce its spread and the agreement with its content (Wojcik et al. 2022). Banas et al. (2022) found that crowdsourced fact-checking is more trusted when it is congruent with another assessment, provided by an AI. However, other research found that providing online users with links and resources from previous users' searches does not improve performance in judging the accuracy of content compared to individuals collecting evidence themselves (Resnick et al. 2021). The present research aims to shed more light on the subject: How is an individual's evaluation of news content influenced by the evaluations of their peers? What factors predict reliance on this information? Do users blindly trust previous ratings, integrate them into their considerations, or disregard them? Do users search for information differently, or take more or less time to evaluate when peer ratings are available?

To address these questions, we devised a preregistered online experiment (osf.io/egkxy) in which participants were asked to rate the scientific validity of a selection of Facebook posts (“How scientifically valid would you rate the information contained in the post?”). We focused on scientific topics because: they are based on empirical facts and thus can be evaluated more objectively than other types of disinformation; they often require technical knowledge to be understood – knowledge that lay people often do not have – making them more opaque and thus particularly vulnerable to misinformation (Branch et al. 2022; Martini and Andreoletti, 2021; Oreskes and Conway, 2010) and individuals posing as experts (Nogara et al. 2022); questions on such topics required a certain amount of effort to answer them; lastly, this type of disinformation often has very direct nefarious consequences on the population, as it concerns fields like health and climate (Suarez-Lledo and Alvarez-Galvez, 2021). The posts included scientific and pseudoscientific content on topics ranging from health to climate change. Alongside each post, we presented participants with a graphical summary of the validity ratings of a group of previous participants. The summary, in the form of a histogram, showed the distribution of these ratings so that participants could understand whether there was a clear consensus or a more dispersed opinion.

We tested whether participants' evaluations of these (pseudo) scientific posts differed from those of a control group (preregistered hypothesis 1). Specifically, we wanted to measure whether previous evaluations that indicated the correct rating increased accuracy, and those that gave a misleading recommendation decreased accuracy. In order to also explore which factors predict the use of previous raters' information, we additionally tested how closely participants imitated the majority based on the shape of the distribution of previous ratings and the personal relevance of the topic (Winter et al. 2015, preregistered

hypothesis 2). Finally, we tested whether the presentation of information from previous raters led to faster response times (preregistered hypothesis 3) and to a change in the use of external sources such as search engines (preregistered hypothesis 4).

Methods and materials

Participants. We recruited 1001 U.K. residents through the online platform *prolific.co* between the 22nd and 23rd of July, 2022. The sample size matches that of the original sample from which the graphical summaries were produced, so as to obtain a balanced statistical comparison (Ronzani et al. 2023). Power analyses based on the achieved effect size and simulations of the main test (interaction between majority opinion and experimental group) yielded an average statistical power of 93.8%. One participant was excluded due to preregistered criteria (using a mobile phone when desktop devices were mandatory, in concordance with previous studies), leading to a final sample of 1000 participants included in the analyses. Average age was 36 ($SD = 13$, three not specified), 50% of participants were female and 58.5% had a Bachelor's degree or higher. The convenience sample was thus younger and more educated than the UK average. Additional information about the sample is described in S5 Analyses.

Design. The experimental design followed the structure of Panizza et al. (2022): during the experiment, participants observed and were able to interact with one out of ten science-themed Facebook posts (Figure 1). Participants' task was to rate the scientific validity of the statements reported in the title, subtitle, and caption of the post ("How scientifically valid would you rate the information contained in the post?"; 6-point Likert scale from (1) "definitely invalid" to (6) "definitely valid"). The 6-point scale allowed to force people to pick a side while communicating on their intuition, while also allowing them to communicate on their uncertainty.

Below the Facebook post reporting the information to be evaluated, participants were presented with a distribution of evaluations made by participants from a previous experiment in Ronzani et al. (2023). Evaluation was self-paced, and participants could leave the page to search for information online. After the evaluation, a questionnaire followed with a series of control questions. Median completion time of the experiment was 5 minutes, and participants were paid £0.70 for their time (around £8.40/hour). To assess the effect of the additional information provided, we compared participants' responses with two groups from the previous experiment that acted as controls (main text and S2 Analyses; see also the section Adherence to preregistration).

Stimuli. Each participant observed one out of ten possible Facebook posts (see S1 File for a full list). Posts varied in terms of: (i) scientific validity of the content (five valid and five invalid posts, either with verified or debunked information) and (ii) topic (six on climate change, four on health and nutrition). All posts came from sources relatively unknown to participants, as measured in a preliminary survey, and confirmed by participants' familiarity ratings. We focused on Facebook posts from sources that were largely unfamiliar to participants because in a previous study familiar news sources were a stronger predictor of how people rated a post than the scientific claims made in the post itself (Panizza et al. 2022). Since the aim of this paper was to investigate the effect of other people's opinions on online users' decision-making, we wanted to avoid any influence of the source on the evaluation of content. This enabled us to concentrate on the impact of the crowd-sourced ratings. Note that the appearance of

posts from unfamiliar sources in a user's news feed is not uncommon, especially through sponsored content, a widely used form of advertising on many social media platforms. The presence of content from sources unfamiliar to the user, often not clearly highlighted on platforms such as Facebook or Instagram, combined with the difficulty for users to distinguish this type of content from that of familiar sources (Wineburg and McGrew, 2016), makes this type of content particularly important to study. Valid and non-valid posts were selected to be balanced as much as possible on a series of variables (bias in reporting, text comprehensibility, interestingness, plausibility and familiarity with the content, see S1 Methods and also Ronzani et al. 2023).

Distribution. Each Facebook post was associated with the real distribution of responses from a group of previous participants (see Figure 2 for an example). Original respondents were informed that their answers would be made available to future participants for reference ("inform others" condition in Ronzani et al. 2023). As the distributions were based on actual answers, one key feature of our manipulation was that for two of the posts – based on the same scientific article linking a higher risk of heart attacks to marijuana consumption in young adults – the most selected validity rating visible in the distribution misleadingly pointed to the incorrect answer. This made it possible to study the influence of previous information both when it was misleading and when it was useful. Notice that all participants were debriefed at the end of the experiment to prevent the diffusion of false beliefs.

Presenting ratings in the form of a distribution differs from other approaches studied in the literature, such as providing comments or refuting references to the incriminated post (Resnick et al. 2021; Wojcik et al. 2022). Compared to these strategies, the use of a distribution summary has two main advantages: ratings can be easily attached to any content, rather than only to those that raise doubts or concerns. A more widespread use of ratings could prevent the implied truth effect, i.e. the impression that repeated content without any fact-checking attached is more likely to be true (Pennycook et al. 2020). In addition, information presented graphically could provide valuable clues about the consensus of peers on the topic or whether there are multiple opinions. Therefore, due to its differences from existing approaches, the current experiment explores several as-yet unexplored factors that could potentially be applied to online platforms.

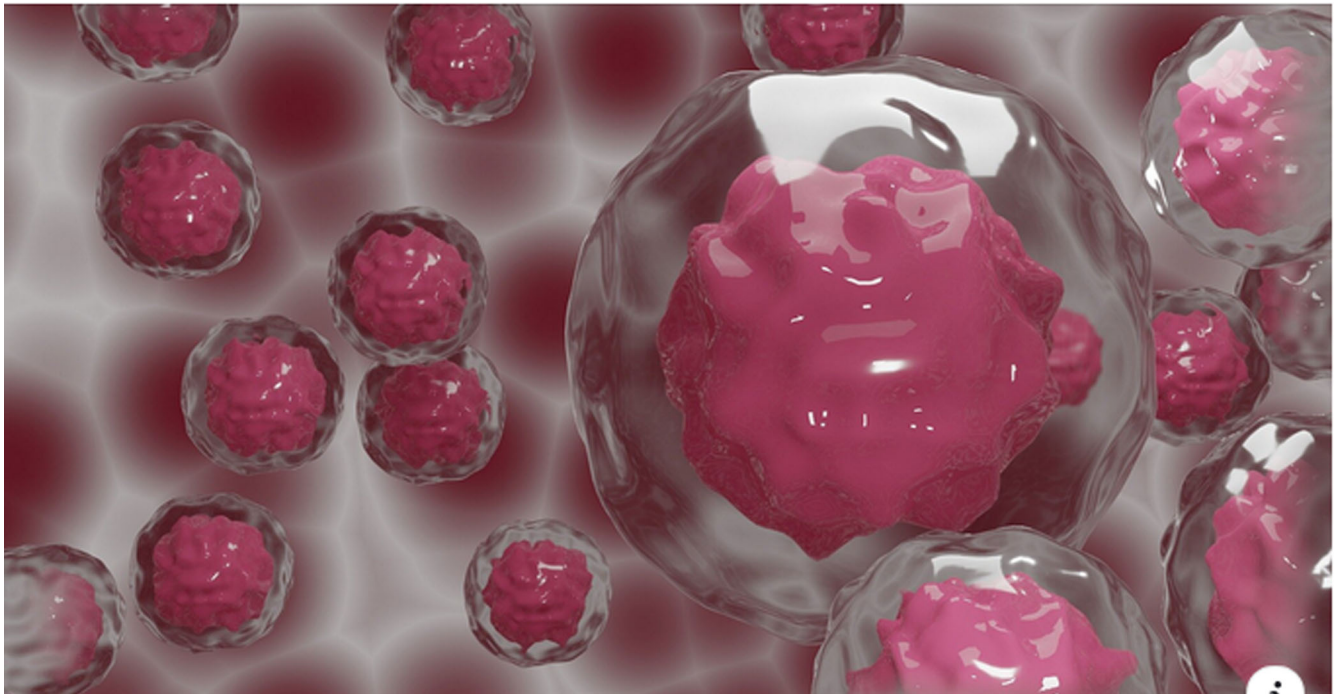
To test whether and how the shape of the distribution affected behaviour, we identified three relevant properties of each distribution (Table 1): a) the majority opinion, i.e., the most selected answer in the distribution, and whether it indicated the correct answer or whether it was misleading; b) peak deviation, i.e., how many participants responded the most selected answer compared to the second-most selected answer ($\#participants_{rst} - \#participants_{second}$), an index of how strong the preference for the most selected answer was, and c) opinion contrast, i.e., whether the second-most selected answer was of opposite sign (e.g., "valid" versus "invalid") compared to the most selected answer, an index of whether or not the distribution pointed unequivocally towards a clear answer.

Hypotheses and predicted variables

Accuracy. The first set of preregistered hypotheses (H1) focused on whether participants who were shown the distribution performed better or worse than the original participants. We computed two measures of performance—accuracy score and correct guessing. Accuracy score is a re-scoring of the validity ratings, ranging from 1 (least accurate) to 6 (most accurate). Specifically,



Climate change will bring an acute toll worldwide, with rising temperatures, wildfires and poor air quality, accompanied by higher rates of cancer, especially lung, skin and gastrointestinal cancers, according to a new report.



MEDICALXPRESS.COM

Climate Change Will Give Rise to More Cancers

New Study Focuses on Global Impact for Major Cancers and Steps Needed to Lesse...



Comments Shares

Like

Comment

Share

Fig. 1 Example of a Facebook post containing valid scientific information. Participants could interact with the post by following its links to the source’s Facebook page, the article, the website, or the info button containing further information such as the link to the source’s Wikipedia page (when present).

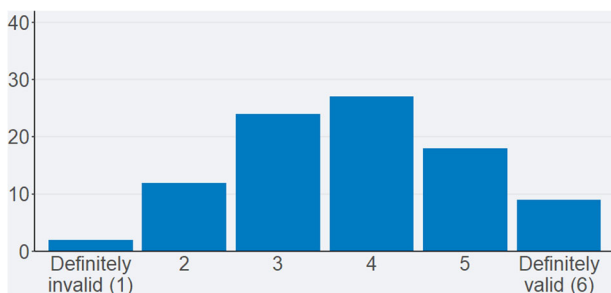


Fig. 2 Example of a response distribution for the Facebook post in Figure R. The post reported findings of a scientifically valid article linking increased rates of several types of cancers to climate change.

an accuracy score of 1 indicates a “definitely invalid (1)” rating when the content of the Facebook post was actually valid, or a “definitely valid (6)” rating when the content of the post was actually invalid. Similarly, an accuracy score of 2 indicates an incorrect “2” or “5” rating, 3 indicates an incorrect “3” or “4” rating, 4 indicates a correct “3” or “4” rating, 5 indicates a correct “2” or “5” rating, and 6 indicates a correct “1” or “6” rating.

The second measure of accuracy, correct guessing, referred instead to a dichotomous variable that simply tracked whether participants gave a ‘valid’ (respectively ‘invalid’) rating when the post content was actually scientifically valid (respectively invalid).

We thus predict that, relative to the original group, the presence of the distribution should: (H1a) increase accuracy scores when the majority opinion is accurate and reduce accuracy scores when the majority opinion is misleading; and (H1b)

Table 1 Properties of each distribution, by post.

Post	Property			
	N	Majority opinion	Peak deviation	Opinion contrast
climate and cancers	92	4 (accurate)	3	True
climate deniers are unqualified	103	4 (accurate)	16	True
end of civilization within 30 years	86	2 (accurate)	14	False
climate science is unsettled	108	2 (accurate)	14	False
equal warming to 1900	96	2 (accurate)	12	False
no climate emergency letter	107	2 (accurate)	3	False
cigars minimal health risks	101	2 (accurate)	4	False
quitting smoking harder for women	97	4 (accurate)	3	True
cannabis and heart attacks	105	3 (misleading)	2	False
cannabis and heart attacks II	107	3 (misleading)	5	False

increase correct guessing when the majority opinion is accurate and reduce correct guessing when the majority opinion is misleading. Both hypotheses are measured by testing the significance of the interaction term between the experimental groups (original group vs. observers) and informativeness of the majority opinion (accurate vs. misleading); additionally, in case of a significant interaction, we conduct post-hoc analyses to test the direction of the results, namely whether the effect is driven by the distributions with the accurate majority opinion, the distributions with the misleading majority opinion, or both.

Opinion-following. The second group of preregistered hypotheses (H2) focused on several predictors of conformity with the majority opinion, including distribution properties and personal relevance of the post’s content (Winter et al. 2015). Would people be sensitive to the prominence of the peak and to contrasting opinions? To test whether and to what extent a participant responded similarly to previous participants, we measured the distance between the participant’s rating and the majority opinion in the original distribution, majority opinion – validity rating, where 0 represents a match between the two.

Our hypotheses regarding opinion following are as follows: (H2a) responses will be closer to the majority opinion when the distribution peak is high, (H2b) responses will be closer to the majority opinion when there is no opinion contrast, and (H2c) responses will be closer to the majority opinion when personal relevance was low. We will test these predictions in the group of observers and test the three predictors (peak deviation, opinion contrast, and personal relevance) simultaneously in a single statistical model.

Response times. Our third set of hypotheses (H3) concerned whether response times were faster compared to the original experimental condition, and whether personal relevance of the content predicted slower response times. The dependent variable was response times, rank-transformed to perform statistical tests.

Specifically, the two hypotheses concerning response times are that (H3a) the observer group should rate the posts more quickly than the original participants, and (H3b) a higher personal relevance of the post’s content will be associated with slower rating times.

Search behaviour. The fourth and last set of hypotheses (H4) tested whether participants made more extensive use of search engines compared to the original group of participants and whether search behaviour was mediated by the personal relevance of content. Use of search engines has in fact been shown to improve the evaluation of unfamiliar content (Panizza et al. 2022). Participants were thus asked whether they engaged in two

search strategies adopted by professional fact-checkers (Breakstone et al. 2019; McGrew et al. 2017): lateral reading (i.e., searching for information on a search engine) and click restraint (i.e., if they further reported looking beyond the first results suggested by the search engine).

Our two hypotheses regarding participants’ search behaviour are that: (H4a) use of lateral reading and click restraint will be reportedly lower in the observer group than in the original group, and (H4b) a higher personal relevance of the post’s content will be associated with higher self-report use of lateral reading and click restraint.

Control measures. In addition to the above measures, we included a series of control measures for our analyses. Other questions included self-report measures of confidence in the validity rating (6-point scale from “don’t know (1)” to “absolutely certain (6)”), intention to share the post (yes/no question), plausibility of the post content (6-point scale from “totally implausible (1)” to “totally plausible (6)”), subjective knowledge about the post’s content (6-point scale from “nothing at all (1)” to “a great deal (6)”), personal relevance of the post’s content (continuous 0-100 rating and a yes/no question), familiarity with the source (yes/no question), perceived trustworthiness of the source (5-point scale from “not at all (1)” to “entirely (5)”), trust in scientists (6-point scale from “not at all (1)” to “a lot (6)”), conspiratorial beliefs (6 5-point scales taken from Bode and Vraga, 2018), altruism (adapted from Rushton et al. 1981), social comparison (adapted from Gibbons and Buunk, 1999), and self-reported strategy to evaluate the post (multiple-choice question; if “I followed previous participants’ answers” is selected, another multiple choice question followed asking why they did so). This last measure was also used to assess whether the participants felt compelled to use the distribution data during their evaluation, a form of experimenter demand. In addition to responses in the questionnaire, we obtained information about participants from the recruiting platform, such as their level of education, socio-economic status, social media use, and belief in climate change.

Analyses. Statistical tests were conducted using R (R Core Team, 2018). We adopted the standard 5% significance level to test against the null hypotheses. All tests were two-tailed unless otherwise specified. Post-hoc tests and multiple comparisons were corrected using the Benjamini-Hochberg procedure, and 95% confidence intervals were also family-wise corrected. Given the small number of stimuli ($N = 10$), we did not cluster errors by post in our regression analyses, although mixed effects analyses—presented in S1 Analyses—yield almost identical results. We follow the specific statistical tests included in the pre-registration unless noted in the section below.

Adherence to preregistration. We conducted our analyses as described in the original preregistration (osf.io/egkxy) with minor changes that we report here. For ease of illustration, participants' performance was compared to that of participants that formed the displayed distributions (inform others condition in Ronzani et al. 2023). The comparison with the control group is presented in S2 Analyses, and results do not deviate meaningfully from those presented in the main text. Two of the original preregistered tests (H3a and H3b) were conducted using linear regression with rank-transformed response times, given that our measure of personal relevance was continuous and the proposed test (Wilcoxon rank-sum test) could not be used with a continuous predictor variable. Two of the exploratory analyses could not be conducted due to statistical issues, one concerning the interaction between experimental condition and personal relevance, since the continuous measure of personal relevance was not present in the original data, and another test including random effects was not possible since predicting variables (distribution properties) were perfectly correlated with random effects. Exploratory analyses that are not presented in the main test are included in S4 and S5 Analyses.

Results

Participants were evenly randomised across posts (Chi-squared test, $\chi^2(9) = .740, p \approx 1$). The median time to evaluate the Facebook post was 36 seconds (minimum: 2 seconds, maximum: 10 minutes), against 34 seconds in the original condition. On a scale from 1 to 6 (3.5 response at chance level), the average accuracy score was 4.07 ($SD = 1.45$), with 64.4% of participants correctly guessing the scientific validity of the post.

Effect of distribution on accuracy. Figure 3 compares the original responses with the ratings given by participants in this experiment. The set of preregistered hypotheses H1 predicted that participants' responses would have differed from the original ones: when the majority opinion was accurate, performance was expected to be higher compared to the original participants, and lower when the majority opinion was misleading. In statistical terms, we expected an interaction between the experimental condition and the informativeness of the distribution (condition \times informative).

This interaction term was significant for both our measures of performance, accuracy score (H1a; ordered logistic regression: $\beta = 0.695[0.314, 1.076], z = 2.405, p < 0.001$) and correct guessing (H1b; logistic regression: $\beta = 0.581[0.109, 1.056], z = 2.405, p = 0.016$). Post-hoc analyses tested whether participants with misleading (respectively correct) information had performed significantly worse (respectively better) than the original participants. The difference in performance was in the predicted direction and significant for accuracy scores (misleading majority opinion: $\beta = -0.413[-0.798, -0.028], z = -2.398, p = 0.017$; accurate majority opinion: $\beta = 0.282[0.081, 0.483], z = 3.136, p = 0.003$), but not significant for correct guessing, perhaps due to smaller statistical power (misleading majority opinion: $\beta = -0.426[-0.905, 0.053], z = -1.988, p = 0.094$; accurate majority opinion: $\beta = 0.155[-0.095, 0.404], z = 1.388, p = 0.165$).

Results were robust to a series of control tests (inclusion of random effects for post, exclusion of participants who reported using the distribution only because this is what the experimenters expected of them, or use of a different control group (S1 and S2 Analyses). The influence of the distribution is even more remarkable considering that 91.5% of the sample reported not having followed previous participants' answers. When the same analyses were repeated including only this subgroup, the

interaction term for accuracy score was still significant ($\beta = 0.732[0.341, 1.123], z = 3.670, p < 0.001$) and marginally significant for correct guessing ($\beta = 0.471[-0.010, 0.955], z = 1.914, p = 0.056$).

The findings concerning accuracy thus fully support H1a and partially support H1b, suggesting that additional information provided to participants with the distribution does indeed affect their evaluation of Facebook posts, although not necessarily while being aware of this influence.

Predictors of following the majority opinion. The second set of hypotheses H2 tested how the properties of the distributions (peak deviation and opinion contrast), combined with the personal relevance of the post's content, predicted how close the participant's response was to the majority opinion. Specifically, we predicted that responses would be closer to previous responses when the distribution peak was high (H2a), when there was no opinion contrast (H2b), and when personal relevance was low (H2c). Indeed, when content is of marginal relevance to users, following the majority might be an efficient low-cost strategy to perform the task.¹

Figure 3 displays a high degree of overlap between participants' responses and majority opinion: the most popular response in the distribution is also the most frequent or second-most frequent response for participants in 9 out of 10 posts. This notwithstanding, opinion-following is not predicted by the properties of the distribution (H2a peak deviation: $\beta = 0.000[-0.031, 0.031], z = 0.001, p = 0.999$; H2b opinion contrast: $\beta = -0.248[-0.612, 0.115], z = -1.633, p = 0.307$) nor by personal relevance (H2c: $\beta = 0.002[-0.005, 0.008], z = 0.589, p = 0.834$).

Lack of significance is illustrated by the post claiming that deniers who oppose climate science are unqualified to make such claims (Figure 2, top right): in this post, the distribution indicated a clear consensus on a '4' response (moderately valid), thus displaying a marked difference between the majority opinion and the second-most rated answer (i.e., a peak deviation of 16 participants). Yet, most participants selected the second-most-rated option, even if less popular by a considerable margin.

Another exploratory analysis tested whether participants' responses clustered around the majority opinion by testing whether the variance in responses was reduced compared to the original variance of the distribution. Contrary to this prediction, however, a Fligner-Killeen test for heterogeneity of variances showed that the variance of responses was actually larger for participants compared to the original distribution ($SD = 1.45$ versus $SD = 1.38$; $\chi^2(1) = 4.662, p = 0.031$). In fact, participants' evaluations were more extreme than the original distribution (ordered logistic regression, $\beta = 0.232[0.068, 0.395], z = 2.780, p = 0.006$).

In summary, neither H2a, H2b nor H2c were supported by the results. On the contrary, it appears that the observer group rated the validity of the contributions more extremely than the original raters.

Response times and search style. Hypotheses H3 and H4 predicted the presence of the distribution would have reduced the time spent on the task and the use of external resources. However, contrary to H3a, participants' response times were not different from those observed in the original condition (median time: 36.3 seconds against 34.4 seconds; linear regression of rank-transformed response times: $\beta = -39[-90, 11], z = -1.525, p = 0.128$). Similarly, contrary to H4a, participants did not make significantly less use of search engines than in the original condition (i.e. no reduced lateral reading, $\beta = -0.009[-0.304, 0.285], z = -0.060, p = 0.952$), nor did

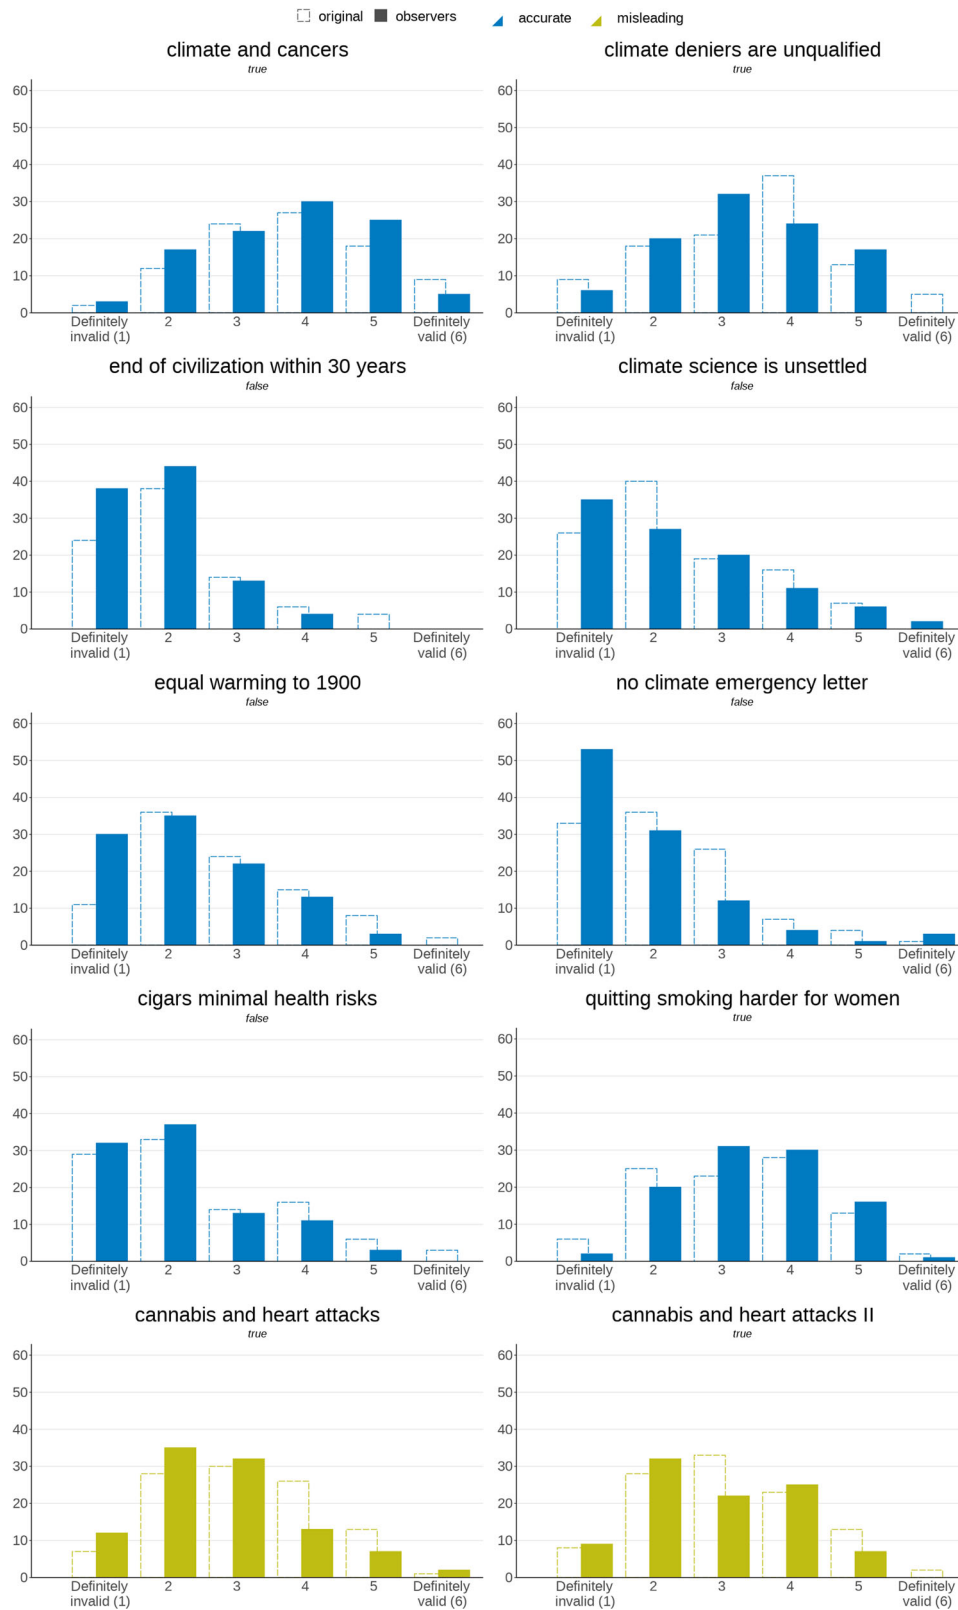


Fig. 3 Comparison of ratings between the original distributions (empty dashed bars) and participants observing the distribution (full bars), by post. Under each post description, it is specified whether the post is valid (true) or not (false). The colour of the bars specifies whether the majority opinion accurately indicated the correct answer or whether it was misleading.

they use it more superficially (i.e. no reduced click restraint, $\beta = -0.047[-0.402, 0.308]$, $z = -0.259$, $p = 0.795$). Additional predictions were that personal relevance should have increased both response times (H3b) and the use of external search (H4b). Self-reported personal relevance however was not a predictor of the duration of the evaluation ($\beta = 1[-1, 2]$, $z = 1.106$, $p = 0.269$) nor of search style (lateral reading: $\beta = 0.006[-0.002, 0.015]$, $z = 1.403$, $p = 0.161$; click restraint: $\beta = 0.009[-0.002, 0.020]$, $z = 1.607$, $p = 0.108$).

Contrary to hypotheses regarding response times and search style, there was no difference between the observer and original groups, nor did personal relevance modulate either variable.

Discussion

Crowdsourcing fact-checking is a solution deemed quite promising in contexts where professional journalists and/or experts' contributions are unavailable or not trusted enough. But to what extent and how do internet users actually rely on their peers' opinions, when they get to evaluate information online?

A first observation from this online experiment is that participants' responses are influenced by the information provided to them about other people's opinions. From a belief-updating perspective, we could indeed expect most people to rationally update their beliefs in light of new information. What is most remarkable is that the vast majority of participants reported not having used previous judgements: in fact, these participants indicated that prior information had no impact on their assessment, while their responses showed that it did. This finding is in line with a third person effect (Davison, 1983), that is the tendency to underestimate the role of cognitive biases on one's own judgements and decisions (as compared to others'). However, ignoring or downplaying the role of external influences on one's own judgements does not necessarily mean that one is not aware of such influences. People might know they are biased, but still want to deny their bias. This could be attributed to several psychological phenomena (e.g., self-presentation, Leary and Kowalski, 1990).

Although we did observe a change in behaviour, blind copying of the majority cannot fully explain the change in performance. Our second observation is indeed that people did not simply copy the most rated option, but rather followed its general trend (that is, whether the post was valid or invalid). It might be that participants used this information as a first indicator of where the truth lies, and then formed a judgement based on that assumption. In other words, the distribution of the previous answers would not be used as a substitute for one's own research, but rather as a prior for one's individual reasoning.

Notably, survey responses suggest that the integration of private assessments and earlier information felt unobtrusive and non-invasive. Such a seamless use of additional information renders this approach a perfect basis for a nudge intervention on social media platforms (Grüning et al. 2023; Lorenz-Spreen et al. 2020; Pennycook and Rand, 2022). An intervention that draws on the wisdom of the crowd may be seen as a form of informative nudge (Acquisti et al. 2017), as it adds valuable information about the observed content without altering the user's freedom of choice. Crucially, information would be presented in a transparent way, without providing summaries, so that users can evaluate the full information (Lorenz-Spreen et al. 2020). Companies may build on this framework or extend their existing programs to report more information than a simple majority rating: even complex information such as the distribution of previous ratings is likely to influence users' decisions, with enough freedom and transparency to allow them to critically evaluate the content under scrutiny.

How participants integrated prior information in our experiment relates to a more general question about the role of social learning (via the observation and imitation of others' attitudes and behaviours; see Bandura and Walters, 1977) in the way humans acquire knowledge online. In social networks in particular, information commonly reaches people indirectly in a two-step process, via their social networks and accounts from their peers (Aarøe and Petersen, 2020). How is such social information taken into account when assessing the relevance and accuracy of online content? While a large number of studies have highlighted the role of popularity and other social cues on information consumption (Nikolov et al. 2019), others have underlined the tendency of individuals to privilege their personal (and sometimes biased) intuitions to the detriment of reliable information from others, especially in the context of conspiracy theories. For instance, the appeal toward "doing one's own research" can be rooted in the desire for autonomy and empowerment and is often fuelled by the distrust of science institutions (Buzzell and Rini, 2023). In fact, it seems that the use of social information is deeply context-dependent, and influenced by the intuitiveness of the information to be evaluated, its implications for oneself, the stakes of knowing the truth of the matter, etc. For example, Acerbi (2019) argues that humans are "weary learners": they sometimes copy others (e.g., either prestigious people or the majority when the latter represents the best source of information available to us), but they do so in a strategic way. For instance, the influence of an advertisement involving a celebrity on people's subsequent behaviour will depend on whether the celebrity is perceived as an expert regarding the information they are communicating about (Amos et al. 2008), and people become more vigilant when the stakes of knowing the truth get higher (Morisseau et al. 2021). In the particular case of this study, people neither had strong prior knowledge regarding the statements they were evaluating nor did they know the persons composing the panel whose responses were provided. Thus, while on the one hand, the trustworthiness of the source was not exceptionally high (the panel was no more expert than the participants themselves), on the other hand, the stakes of knowing the truth of the matter were relatively low too, as participants did not suffer any immediate consequences from believing false information. In such a situation, prior participants' behaviour can be seen as a strategic trade-off between lower cognitive effort as a result of the trust placed in the panel, and an individual search motivated by the context of the experiment, curiosity, the desire to give a good answer, and so on. The ability to make good epistemic trade-offs when assessing information online—both for oneself and for the sake of one's online community, by avoiding the spreading of misinformation—can benefit from such crowd-sourced solutions, especially when the truth of the matter is not of paramount importance to the online user.

Another relevant finding is that neither the time taken to give one's response, nor the use of strategies to evaluate the information are affected overall by the presence of the distribution. This result is at odds with the idea that people would rely solely on the majority opinion to provide their answers. As the time spent and the strategies deployed did not differ from that of the original participants, this might indicate, on the contrary, an unwavering motivation to seek confirmation (or possibly refutation) of the majority trend. Perhaps, the presence of the distribution of other people's answers made the search for information all the more relevant, either towards more detailed information or towards the possible invalidation of the majority answer. However, we remain cautious in making speculations at the current stage, as other factors may have contributed to the current non-significant differences. For instance, as response times are concerned, it should be noted that participants observed

more information on the screen (i.e., text and graphical summary) compared to the original participants (i.e., text only), which might have resulted in longer reading times. Similarly, the fact that the personal relevance of the post's content did not predict response times or search style might be equally affected by factors initially unaccounted for. In our pre-registered hypotheses, relevance should have predicted longer evaluation times and more frequent and more thorough searches. Instead, these relations might not be univocal: higher relevance might translate into higher prior knowledge and therefore faster and more frugal responses. Also the contrary could be true: some people who find the content personally relevant might lack the particular knowledge to assess its accuracy and might choose to follow the majority. Future research could disentangle factual knowledge from personal relevance, as well as control for the amount of information presented.

While the present study brings numerous insights into the reception of peer fact-checking, it also comes with limitations to the generalisation of its findings. First, many of the posts gathered as stimuli were months to years old at the time of data collection, contrary to other studies (Godel et al. 2021). Older articles ensured that the claims presented could be compared against a clear scientific record and thus allowed to properly measure the accuracy of evaluation. Although scientific claims are not likely to become outdated as quickly as other types of news, recent articles present a particular challenge to the reader that is worth investigating, given the scattered and often contradictory information on which they are based. Another limitation is that articles in the experiment were selected by experimenters and thus might not be representative of the content observed in a typical newsfeed (Allen et al. 2021). As a result, posts presented to participants might have been off-target with respect to the user's characteristics and interests. Future studies could resort to more ecologically valid and recent news to investigate whether current observations extend to such content.

Similarly, while our focus was on targeting content from unknown sources, users' newsfeed diets consist overwhelmingly of content from credible sources (Guess et al. 2021). While it could be argued that a smaller proportion of potentially misleading content does not reduce its impact, our results are indeed less generalisable to all instances of online information consumption. Given that online users are less likely to consider additional cues when they have priors about either the source or the information conveyed, we might expect crowdsourced advice to have a lower impact on known sources. Thus, it is crucial for future research to investigate how crowdsourced interventions are received when targeting known sources, and test whether differences in source prevalence across platforms also affect reception.

Another limitation is that the way in which scientific validity was rated (i.e., 6-point scale) may have played a role in the way users used the opinion of previous participants. For instance, under the experimental conditions, expecting a majority of "5s" when "5" is the most selected option by previous participants does not take into account the base-rate probability of selecting "5", which in this specific metric equals 16.6%. In other words, we cannot exclude that the tendency to copy prior participants' judgements could have been amplified if scientific validity was assessed in a "yes" vs. "no" format. Base rates might also explain why neither personal relevance nor the shape of the distribution predicted who responded like the majority and who did not. Given these considerations, future studies should investigate how the rating format might affect the evaluation and integration of ratings.

Our sample was also on average younger and more educated than social media users in the U.K. (Statista, 2023), which might raise the question of how an older, less educated sample would respond to ratings of previous users. Previous research on

demographic factors influencing the spread of misinformation suggests that education level does not significantly affect the likelihood of sharing misinformation (Guess et al. 2019), while age may increase sharing (Guess et al. 2019) or have a null effect (Pehlivanoglu et al. 2022). Based on these findings, we speculate that including older or less educated participants in our study would not reduce the prevalence of misinformation beliefs, and may even increase it, creating more room for other users' ratings to have an impact. This prediction however will need to be tested in future studies investigating potential heterogeneity across online users.

While our study focused on scientific information and misinformation, the present approach should be tested on a variety of other claims, such as those focusing on political facts, economic figures, historical events, and social issues. We expect that the approach we studied should have similar results in other domains, and that it should be most effective for claims that are new to users and around which they have not engaged in motivated reasoning (Druckman and McGrath, 2019). These hypotheses are left to be tested in future studies that aim to extend the scope of the present research.

Another important limitation is the nature of the crowdsourced evaluation: while participants in the experiment observed the distribution of opinions from previous participants, other formats have been used on social media platforms. Twitter's Community Notes program, for instance, allows evaluators to plot a summary of the checked claims and add updated references to provide context. Despite the differences, our current results provide insights also for this type of fact-checking: users will potentially benefit from the presence of additional information, but will likely continue to double-check claims on their own. One potential avenue for expanding the current research will be to test whether an established community, which has built a reputation for providing accurate fact-checking, can induce greater trust in its assessments.

Conclusions

In this study, we aimed to test whether informing about the evaluation of a given content from peers influences one's evaluation of this content, and to identify the factors that predict this influence. We found that, despite being influenced by the evaluation of previous participants, responses follow a more nuanced pattern than the mere copying of the most rated option. Our results suggest that individual reasoning plays a significant role when one assesses information online, and that crowd-sourced information can help internet users be more resilient to misinformation.

Data availability

All data are available on the osf repository (osf.io/c6qka).

Received: 11 April 2023; Accepted: 27 October 2023;

Published online: 25 November 2023

Note

1 Although relations between relevance and response strategies might not be univocal, a point that we address in the discussion.

References

- Aarøe L, Petersen MB (2020) Cognitive biases and communication strength in social networks: The case of episodic frames. *Br J Pol Sci* 50:1561–1581
- Acerbi A. (2019) Cultural evolution in the digital age. Oxford University Press
- Acerbi A, Altay S, Mercier, H (2022) Research note: Fighting misinformation or fighting for information? *Harvard Kennedy School (HKS) Misinform Rev*

- Acquisti A et al. (2017) Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Comput Surv (CSUR)* 50:1–41
- Allen J, Arechar AA, Pennycook G, Rand DG (2021) Scaling up fact-checking using the wisdom of crowds. *Sci Adv* 7:eabf4393
- Altay S, Acerbi A (2023) People believe misinformation is a threat because they assume others are gullible. *New Med Soc* 14614448231153379
- Álvarez-Moreno C, Valderrama-Beltrán S, Rodríguez-Morales AJ (2021) Implications of antibiotic use during the COVID-19 pandemic: the example of associated antimicrobial resistance in Latin America. *Antibiotics* 10(3):328
- Amos C, Holmes G, Stratton D (2008) Exploring the relationship between celebrity endorser effects and advertising effectiveness: A quantitative synthesis of effect size. *Int J Adv* 27:209–234
- Banas JA et al. (2022) When machine and bandwagon heuristics compete: Understanding users' response to conflicting ai and crowdsourced fact-checking. *Human Commun Res* 48:430–461
- Bandura A, Walters RH (1977) *Social learning theory*, vol. 1. Englewood cliffs Prentice Hall
- Bode L, Vraga EK (2018) See something, say something: Correction of global health misinformation on social media. *Health Commun* 33:1131–1140
- Branch T, Origgi G, Morisseau T (2022) Why trust raout? how social indicators inform the reputations of experts. *Soc Epistemol* 36:299–316
- Breakstone J. et al. (2019) Students' civic online reasoning: A national portrait. *Educ Res* 0013189X211017495 (2019)
- Breitsohl J, Wilcox-Jones JP, Harris I (2015) Groupthink 2.0: An empirical analysis of customers' conformity-seeking in online communities. *J Customer Beh* 14:87–106
- Buzzell A, Rini R (2023) Doing your own research and other impossible acts of epistemic superheroism. *Phil Psychol* 36:906–930
- Coan TG, Boussalis C, Cook J, Nanko MO (2021) Computer-assisted classification of contrarian claims about climate change. *Sci Rep* 11:1–12
- Coleman J (2021) Introducing birdwatch, a community-based approach to misinformation. https://blog.twitter.com/en_us/topics/product/2021/introducing-birdwatch-a-community-based-approach-to-misinformation
- Colliander J (2019) "this is fake news": Investigating the role of conformity to other users' views when commenting on and spreading disinformation in social media. *Comp Hum Behav* 97:202–215
- Cook J, Ecker U, Lewandowsky S (2015) Misinformation and how to correct it. In: Scott RA, & Kosslyn SM (eds), *Emerging Trends in the Social and Behavioral Sciences: An Interdisciplinary, Searchable, and Linkable Resource*. John Wiley & Sons, p 1–17. <https://doi.org/10.1002/9781118900772.etrds0222>
- Davison WP (1983) The third-person effect in communication. *Publ Opinion Quar* 47:1–15
- Druckman JN, Levendusky MS, McLain A (2018) No need to watch: How the effects of partisan media can spread via interpersonal discussions. *Am J Pol Sci* 62:99–112
- Druckman JN, McGrath MC (2019) The evidence for motivated reasoning in climate change preference formation. *Nat Climate Change* 9:111–119
- Fazio L (2020) Pausing to consider why a headline is true or false can help reduce the sharing of false news. *Harvard Kennedy School Misinform Rev* 1
- Gibbons FX, Buunk BP (1999) Individual differences in social comparison: development of a scale of social comparison orientation. *J Pers Social Psychol* 76:129
- Godel W. et al. (2021) Moderating with the mob: Evaluating the efficacy of real-time crowdsourced fact-checking. *J Online Trust Saf*. 1
- Grüning DJ, Riedel F, Lorenz-Spreen P (2023) Directing smartphone use through the self-nudge app one sec. *Proc Nat Acad Sci* 120:e2213114120
- Guess A, Aslett K, Tucker J, Bonneau R, Nagler J (2021) Cracking open the news feed: Exploring what us facebook users see and share with large-scale platform data. *J Quant Descr: Dig Media* 1
- Guess A, Nagler J, Tucker J (2019) Less than you think: Prevalence and predictors of fake news dissemination on facebook. *Science Adv* 5:eaa4586
- Hertwig R, Grüne-Yanoff T (2017) Nudging and boosting: Steering or empowering good decisions. *Persp Psychol Sci* 12:973–986
- Jang SM, Kim JK (2018) Third person effects of fake news: Fake news regulation and media literacy interventions. *Comput Hum Beh* 80:295–302
- Leary MR, Kowalski RM (1990) Impression management: A literature review and two-component model. *Psychol Bull* 107:34
- Lewandowsky S (2021) Climate change, disinformation, and how to combat it. *Ann Rev Public Health*
- Lewandowsky S, Van Der Linden S (2021) Countering misinformation and fake news through inoculation and prebunking. *Eur Rev Soc Psychol* 1–38
- Lorenz-Spreen P et al. (2021) Boosting people's ability to detect microtargeted advertising. *Sci Rep* 11:1–9
- Lorenz-Spreen P, Lewandowsky S, Sunstein CR, Hertwig R (2020) How behavioural sciences can promote truth, autonomy and democratic discourse online. *Nat Hum Behav* 1–8
- Martel C, Allen J, Pennycook G, Rand DG (2023) Crowds can effectively identify misinformation at scale. *Persp Psychol Sci* 17456916231190388
- Martini C, Andreoletti M (2021) Genuine versus bogus scientific controversies: The case of statins. *History Philosophy Sci* 43:1–23
- McGrew S, Ortega T, Breakstone J, Wineburg S (2017) The challenge that's bigger than fake news: Civic reasoning in a social media environment. *Am Educ* 41:4
- Morisseau T, Branch T, Origgi G (2021) Stakes of knowing the truth: a motivational perspective on the popularity of a controversial scientific theory. *Front Psychol* 3800
- Nikolov D, Lalmas M, Flammini A, Menczer F (2019) Quantifying biases in online information exposure. *J Assoc Inf Sci Technol* 70:218–229
- Nogara G. et al. (2022) The disinformation dozen: An exploratory analysis of covid-19 disinformation proliferation on twitter. In *Proceedings of the 14th ACM Web Science Conference 2022*, WebSci '22, 348–358 (Association for Computing Machinery, New York, NY, USA). <https://doi.org/10.1145/3501247.3531573>
- Oreskes N, Conway EM (2010) Defeating the merchants of doubt. *Nature* 465:686–687
- Panizza F et al. (2022) Lateral reading and monetary incentives to spot disinformation about science. *Sci Rep* 12:1–15
- Pehlivanoglu D, Lighthall NR, Lin T, Chi KJ, Polk R, Perez E, Cahill BS, Ebner NC (2022) Aging in an "infodemic": The role of analytical reasoning, affect, and news consumption frequency on news veracity detection. *J Exp Psychol Appl* 28(3):468–485. <https://doi.org/10.1037/xap0000426>
- Pennycook G, Bear A, Collins ET, Rand DG (2020) The implied truth effect: Attaching warnings to a subset of fake news headlines increases perceived accuracy of headlines without warnings. *Manag Sci* 66:4944–4957
- Pennycook G et al. (2021) Shifting attention to accuracy can reduce misinformation online. *Nature* 1–6
- Pennycook G, Rand DG (2022) Nudging social media toward accuracy. *Annals Am Acad Polit Soc Sci* 700:152–164
- R Core Team. (2018) *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>
- Resnick P, Alfayez A, Im J, Gilbert E (2021) Informed crowds can effectively identify misinformation. *arXiv preprint arXiv:2108.07898* (2021)
- Ronzani P, Panizza F, Mattavelli S, Morisseau T, Martini C. (2023) How different incentives reduce scientific misinformation online. *Harvard Kennedy School Misinform Rev*
- Roozenbeek J, Culloty E, Suiter J (2023) Countering misinformation. *Eur Psychol* 28(3):189–205
- Rossini P, Stromer-Galley J, Baptista EA, Veiga de Oliveira V (2021) Dysfunctional information sharing on whatsapp and facebook: The role of political talk, cross-cutting exposure and social corrections. *New Media Soc* 23:2430–2451
- Rushton JP, Chrisjohn RD, Fekken GC (1981) The altruistic personality and the self-report altruism scale. *Pers Individual Diff* 2:293–302
- Statista (2023) Social media users in the united kingdom (uk) in 2023. <https://www.statista.com/statistics/1401035/uk-social-media-users-by-education/>
- Ștefăniță O, Corbu N, Buturoiu R (2018) Fake news and the third-person effect: They are more influenced than me and you. *J Media Res* 11
- Suarez-Lledo V, Alvarez-Galvez J (2021) Prevalence of health misinformation on social media: systematic review. *J Med Int Res* 23:e17187
- Sundar SS (2008) *The MAIN model: A heuristic approach to understanding technology effects on credibility*. MacArthur Foundation Digital Media and Learning Initiative Cambridge, MA
- Sunstein CR (2016) *The ethics of influence: Government in the age of behavioral science*. Cambridge University Press
- Tandoc EC (2021) Fake news. In *The Routledge Companion to Media Disinformation and Populism*. Routledge, 110–117
- Thaler RH, Sunstein CR (2019) *Nudge: Improving decisions about health, wealth, and happiness*. Penguin
- Tsikerdekis M (2013) The effects of perceived anonymity and anonymity states on conformity and groupthink in online communities: Aw wikipedia study. *J Am Soc Inf Sci Technol* 64:1001–1015
- Tully M, Maksl A, Ashley S, Vraga EK, Craft S. Defining and conceptualizing news literacy. *Journalism* 14648849211005888 (2021)
- Wineburg S, McGrew S. (2016) Evaluating information: The cornerstone of civic online reasoning. *Tech. Rep.*, Stanford History Education Group, Robert R. McCormick Foundation (2016)
- Winter S, Brückner C, Krämer NC (2015) They came, they liked, they commented: Social influence on facebook news channels. *Cyberpsychol, Beh Social Netw* 18:431–436
- Wojcik S. et al. Birdwatch: Crowd wisdom and bridging algorithms can inform understanding and reduce the spread of misinformation. *arXiv preprint arXiv:2210.15723* (2022)
- Yoo J, Kim D, Kim W-G (2022) Fake news on you, not me: the third-person effects of fake news in south korea. *Commun Res Rep* 39:115–125
- Zhu H, Huberman B, Luon Y (2012) To switch or not to switch: understanding social influence in online choices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2257–2266

Zuckerberg M, Zittrain J (2019) Facebook <https://www.facebook.com/zuck/videos/10106612617413491/>

Acknowledgements

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 870883. The information and opinions are those of the authors and do not necessarily reflect the opinion of the European Commission.

Author contributions

FP conceived the idea, FP, PR, TM, SM and CM developed the idea, FP built the experiment, FP conducted the experiment, FP wrote the original draft, PR, TM, SM and CM revised the manuscript, CM secured the funding.

Competing interests

The authors declare no competing interest.

Ethical approval

The experimental protocols were approved by the Research Ethics Committee (CER) at the University Paris Cite (IRB No: 00012021-05), and all research was performed in accordance with the relevant guidelines and regulations.

Informed consent

All participants gave their written informed consent for participating in the experiment.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1057/s41599-023-02329-y>.

Correspondence and requests for materials should be addressed to Folco Panizza.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023