

MIT Open Access Articles

Emphasizing publishers does not effectively reduce susceptibility to misinformation on social media

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

Citation: Dias, Nicholas, Pennycook, Gordon and Rand, David G. 2020. "Emphasizing publishers does not effectively reduce susceptibility to misinformation on social media." Harvard Kennedy School Misinformation Review.

As Published: 10.37016/MR-2020-001

Publisher: Shorenstein Center for Media, Politics, and Public Policy

Persistent URL: <https://hdl.handle.net/1721.1/144236>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution 4.0 International license





Research Article

Emphasizing publishers does not effectively reduce susceptibility to misinformation on social media

Survey experiments with nearly 7,000 Americans suggest that increasing the visibility of publishers is an ineffective, and perhaps even counterproductive, way to address misinformation on social media. Our findings underscore the importance of social media platforms and civil society organizations evaluating interventions experimentally, rather than implementing them based on intuitive appeal.

Authors: Nicholas Dias (1), Gordon Pennycook (2), David G. Rand (3)

Affiliations: (1) Communication and Political Science, Annenberg School for Communication and the School of Arts & Sciences Gerald; (2) Behavioral Science at University of Regina's Hill/Levene Schools of Business; MIT Sloan (3)

How to cite: Dias, Nicholas; Pennycook, Gordon; Rand; David G. (2020). Emphasizing publishers does not effectively reduce susceptibility to misinformation on social media, *The Harvard Kennedy School (HKS) Misinformation Review*, Volume 1, Issue 1
Received: Oct. 31, 2019 Accepted: Dec. 5, 2019 Published: Jan. 14, 2020

Research question

Platforms are making it easier for users to identify the sources (i.e., the publisher) of online news. Does this intervention help users distinguish between accurate and inaccurate content?

Essay summary

- In two survey experiments, we showed American participants (N=2,217) a series of actual headlines from social media, presented in Facebook format.
- For some headlines, we made publisher information more visible (by adding a logo banner), while for others we made publisher information less visible (by removing all publisher information).
- We found that publisher information had no significant effects on whether participants perceived the headline as accurate, or expressed an intent to share it – regardless of whether the headline was true or false. In other words, seeing that a headline came from a misinformation website did not make it less believable, and seeing that a headline came from a mainstream website did not make it more believable.
- To investigate this lack of an effect, we conducted three follow-up surveys (total N=2,770). This time, we asked participants to rate either (i) the trustworthiness of a range of publishers or (ii) the accuracy of headlines from those publishers, given no source information (i.e., the headline's "plausibility").
- We found a strong, positive correlation between trust in a given outlet and the plausibility of headlines it published. Thus, in many cases, learning the source of a headline does not appear to

¹ A publication of the Shorenstein Center on Media, Politics and Public Policy at Harvard University's John F. Kennedy School of Government.

add information about its accuracy beyond what was immediately apparent from the headline alone.

- Finally, in a survey experiment with 2,007 Americans, we found that providing publisher information only influenced headline accuracy ratings when headline plausibility and publisher trust were “mismatched” – for example, when a headline was plausible but came from a distrusted publisher (e.g., fake-news or hyperpartisan websites).
- In these cases of mismatch, identifying the publisher reduced accuracy ratings of plausible headlines from distrusted publishers, and increased accuracy ratings of implausible headlines from trusted publishers.
- However, when we fact-checked the 30% of headlines from distrusted sources in our set that were that were rated as plausible by participants, we found they were mostly true. In other words, providing publisher information would have increased the chance that these true headlines would be mistakenly seen as false – raising the possibility of unintended negative consequences from emphasizing sources.
- Our results suggest that approaches to countering misinformation based on emphasizing source credibility may not be very effective and could even be counterproductive in some circumstances.

Implications

Scholars are increasingly concerned about the impact of misinformation and hyperpartisan content on platforms such as Facebook and Twitter (Lazer et al., 2018). High-profile criticisms from researchers, journalists and politicians – and the threat of regulation – have pressured social media companies to engineer remedies. Reluctant to be arbiters of truth, these companies have emphasized providing contextual information to users so they can better judge content accuracy themselves (Hughes et al., 2018; Leathern, 2018; Pennycook, Bear, Collins, and Rand, in press; Samek, 2018). Much of that contextual information concerns the sources of content. For example, Facebook’s ‘Article Context’ feature surfaces encyclopedic entries for the sources of articles linked in posts (Hughes et al., 2018). Similarly, YouTube ‘notices’ signal to users when they are consuming content from government-funded organizations (Samek, 2018).

Neither Facebook nor YouTube has released data about the effectiveness of their source-based interventions, and the existing academic literature is inconclusive. Many studies have shown that readers are more likely to believe messages from high-credibility sources relative to identical messages from low-credibility sources (Pornpitakpan, 2004). However, these effects have varied from study to study, likely due to changes in the type of content under review and the characteristics of the sources involved. Moreover, some researchers have argued that seemingly trivial differences in communication media can significantly change how people process messages (Sundar, 2008). Studies examining whether differences in news organization credibility can affect perceptions of news veracity have been similarly inconclusive. Many studies have found that source credibility can affect judgments of accuracy or believability (Baum & Groeling, 2009; Landrum et al., 2017; Berinsky, 2017; Swire et al., 2017; Knight Foundation, 2018; Kim et al., 2019). Yet, others (Austin & Dong, 1994, Pennycook & Rand, 2019b, Jakesch et al., 2018) have failed to demonstrate that assessments of news stories are changed by source information.

The research we present here confirms previous findings that U.S. adults, on average, tend to be good at identifying which sources are not credible (Pennycook & Rand, 2019a). Nonetheless, we find that, for most headlines, emphasizing publisher information has no meaningful impact on evaluations of headline accuracy. Thus, our findings suggest that emphasizing sources as a way to counter misinformation on social media may be misguided. Moreover, our analysis of plausible headlines from fake and hyperpartisan news sites suggests that source-based interventions of this kind may even be counterproductive, as

plausible headlines from distrusted sources are often truthful. These observations underscore the importance of social media platforms and civil society organizations rigorously assessing the impacts of interventions (source-based and otherwise), rather than implement them based on intuitive appeal

Findings

Finding 1: Emphasizing publishers did not affect the perceived accuracy of either false or true headlines.

In Study 1, 562 Americans from the online labor market Amazon Mechanical Turk (MTurk; Horton et al., 2011) were shown 24 headlines in the format of the Facebook News Feed (Figure 1). Half were true, half were false. Half were sympathetic to Democrats, half were sympathetic to Republicans – creating four groups of headlines. Participants were randomly assigned to one of three conditions that varied the visibility of the headlines’ publishers and asked the participants to judge the accuracy of, and whether they would share, each headline on social media.




Figure 1. Source visibility manipulation example. Sample false (left) and true (right) headlines as shown to participants in the logo-banner condition of Studies 1 and 2. In the no-source condition, neither the grey-text domain nor the logo banner was shown. In the Facebook baseline condition, only the grey-text domain was shown.

Surprisingly, we found that source visibility had no significant effect on participants’ media truth discernment for accuracy judgments (Figure 2a) or sharing intentions (Figure 2b): the difference between true and false headlines was equivalent regardless of whether information about the headline’s publisher was absent, indicated in grey text (as on Facebook), or highlighted with a large logo banner (interaction between condition and headline veracity: accuracy judgments, $F=1.59$, $p=.20$; sharing intentions, $F=2.68$, $p=.068$).




Figure 2. Varying source visibility did not affect accuracy perceptions or sharing intentions. Average reported accuracy (top row) and probability of sharing on social media (bottom row) for false versus true headlines across source visibility conditions in Study 1 (left columns) and Study 2 (right columns). Error bars indicate 95% confidence intervals.

In Study 2, we sought to test whether the null effects of adding the logo banner in Study 1 were merely the result of insufficient statistical power. Thus, we replicated the Facebook-baseline and logo-banner conditions from Study 1 using a substantially larger sample ($N=1,845$ Americans from MTurk). As shown in Figures 2c and 2d, the results of Study 2 confirm the ineffectiveness of emphasizing publishers on accuracy judgments and sharing intentions (interaction between condition and veracity: accuracy, $F=0.02$, $p=.89$; sharing, $F=0.00$, $p=.97$).

Finally, we pooled the Study 1 and 2 data from the Facebook-baseline and logo-banner conditions for maximal power and investigated whether the effect of the banner differed based on (i) the participant's partisanship, proxied by whether the participant preferred Donald Trump or Hillary Clinton for president, and (ii) whether the headline was favorable or unfavorable towards the participant's party. We found no evidence that the banner increased accuracy or sharing discernment, regardless of participant partisanship or headline slant. We also found no evidence that the effect of the banner differed among participants with versus without a college degree, or that our results changed when only considering ideologically extreme participants. Finally, focusing only on the one (true) headline from foxnews.com (which Trump supporters trusted strongly while Clinton supporters distrusted strongly), we continue to find no effect of the banner. Statistical details for all analyses are shown in the Supplementary Materials.

Finding 2: There is a strong positive correlation between (i) how trusted a given outlet is and (ii) how plausible headlines from that publisher seem when presented with no source information.

Why did emphasizing sources not improve truth discernment in Studies 1 and 2? One explanation, which we advance here, could be that publisher information often confirms what is directly observable from the headlines. That is, emphasizing the publisher should only change accuracy ratings if one's trust in the publisher is mismatched with the plausibility of the headline in the absence of source information (see Figure 3a). For example, if an implausible headline is published by a trusted news source, emphasizing the publisher should *increase* the perceived accuracy of the headline. Conversely, if a plausible headline is published by a distrusted source, emphasizing the publisher should *reduce* the perceived accuracy of the headline. However, when trusted sources publish plausible headlines, and distrusted sources publish implausible headlines, information about the publisher offers little beyond what was apparent from the headline. Thus, emphasizing the publisher should have no meaningful effect. We found support for this explanation in Studies 3 - 5 by correlating trust in a given source with the plausibility of headlines from that source.

In Study 3, we asked 250 Americans from MTurk to indicate how much they trusted the sources that published the 24 headlines used in Studies 1 and 2. Consistent with our hypothesis, we found a large positive correlation, $r=.73$, between the perceived accuracy of each headline in the no-source condition of Study 1 and trust in the headline's source measured in Study 3 (Figure 3b). Given this alignment, it is not surprising that source visibility had little impact in Studies 1 and 2. But do these 24 headlines represent the more general relationship between source credibility and headline plausibility observed in the contemporary American media ecosystem? We address this question in Study 4 and provide a replication using a more demographically representative sample in Study 5. To do so, we examined a set of 60 sources that equally represented mainstream, hyperpartisan, and fake-news publishers and included some of the most shared sites in each category (Pennycook & Rand, 2019a). From each source, we selected the 10 best-performing headlines that made claims of fact or opinion and were published between August 2017 and August 2018, creating a representative corpus of 600 headlines.




Figure 3. Relationship between trust in outlet and headline plausibility. Diagram of expected impact of emphasizing source given (mis)match outlet trust and headline plausibility when no information was provided (a) and data from Studies 3 (b), 4 (c), and 5(d). Each data point represents one headline. In panels b-d, 0 corresponds to the minimum value of the trust scale and 1 corresponds to the maximum value.

We then assigned participants from MTurk and the sampling company Lucid (Coppock & McClellan, 2019) to either rate the trustworthiness of a random subset of the 60 sources (N=251 from MTurk in Study 4, N=253 from Lucid in Study 5) or rate the accuracy of a random subset of the 600 headlines (N=1,008 from MTurk in Study 4, N=1,008 from Lucid in Study 5). The results are shown in Figure 3c and 3d. Consistent with prior work (e.g. Pennycook & Rand 2019a,b), non-preregistered tests showed that headlines from mainstream sources were rated as more accurate than headlines from hyperpartisan or fake-news sources, and that mainstream sources were trusted more than hyperpartisan or fake-news sources ($p < .001$ for all; see Supplement for details).

Critically, we observed a strong, positive correlation between the plausibility of a headline when no publisher information was provided and trust in the headline's publisher (Study 4, $r = 0.53$; Study 5, $r = 0.43$). For most headlines, the accuracy rating matched the publisher's trustworthiness, such that increasing the

source's visibility should not affect accuracy ratings (as observed in Studies 1 and 2). This pattern was robust when considering Democrats and Republicans separately (see Supplementary Materials). Interestingly, to the extent that headline accuracy ratings did deviate from source trust ratings, they did so in an asymmetric way: there were few implausible headlines from trusted sources, but a fair number of plausible headlines from distrusted sources. Due to this mismatch, it is possible that emphasizing the sources of this latter category of headlines could decrease their perceived accuracy. We tested this possibility in Study 6.

Finding 3: Identifying a headline's publisher only affects accuracy judgments inasmuch as there is a mismatch between the headline's plausibility and the publisher's trustworthiness.

In Study 6, we had 2,007 Americans from MTurk rate the accuracy of 30 headlines (presented without a lede or image). Participants were randomly assigned to either see no information about the publisher (control condition) or see boldface text beneath the headline indicating the domain that published the headline. The 30 headlines were selected at random from a subset of 252 headlines from those used in Studies 4 and 5 (see Methods for details of headline selection). Headlines varied in both their accuracy ratings and whether they were published by distrusted (hyperpartisan or fake-news) or trusted (mainstream) sites.

As per our mismatch account, we found that the impact of naming a headline's publisher varied by the trustworthiness of the publisher and the plausibility of the headline (as measured by accuracy ratings in the control condition). For headlines from hyperpartisan or fake-news publishers (Figure 4a), identifying the publisher had little effect on relatively implausible headlines, but reduced the perceived accuracy of plausible headlines (correlation between headline plausibility and publisher identification effect $r = -0.415$, bootstrapped 95% CI [-.504, -.325]). Conversely, for headlines from mainstream news publishers (Figure 4b), identifying the publisher had little effect on relatively plausible headlines, but increased the perceived accuracy of implausible headlines (correlation between headline plausibility and publisher identification effect $r = -0.371$, bootstrapped 95% CI [-.505, -.294]). See Supplementary Materials for statistical details. These results provide empirical support for the hypothesis we outlined in Figure 3a.

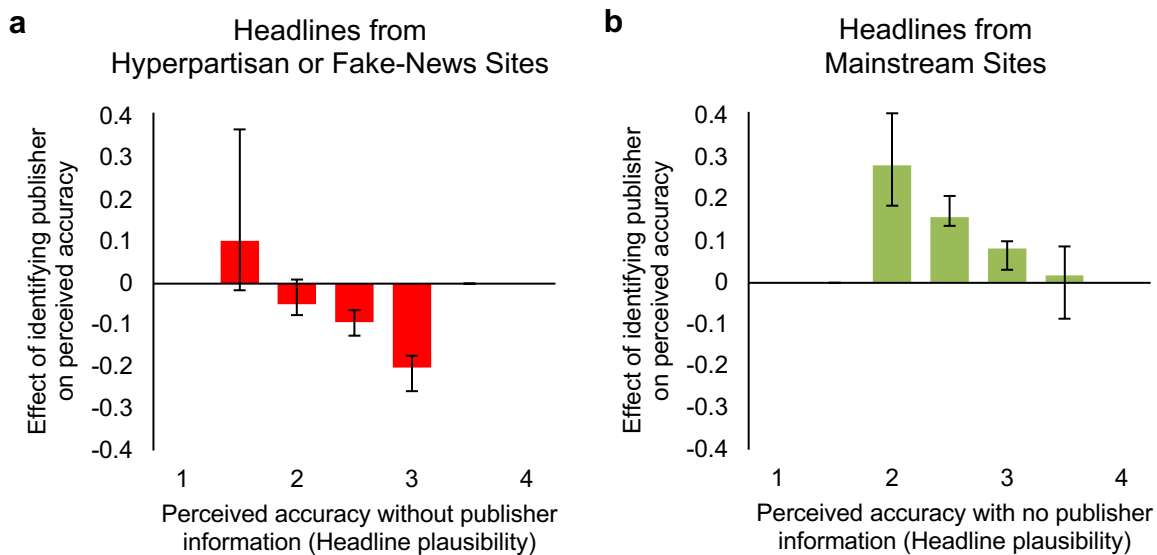


Figure 4. Identifying a headline's publisher only changes accuracy perceptions inasmuch as source trust and headline plausibility are mismatched. Headline plausibility as rated in the no-source control condition (x-axis) plotted against the change

in accuracy judgments when identifying the headline's publisher (y-axis), for distrusted sources (a) and trusted sources (b). Error bars indicate bootstrapped 95% confidence intervals.

Finding 4: Most plausible headlines from distrusted sources were actually true, such that decreasing their perceived accuracy would reduce truth discernment.

Finally, to assess whether emphasizing the source of mismatched headlines has the potential to improve truth discernment, we fact-checked the 120 headlines from distrusted sources that participants in both Studies 4 and 5 rated as accurate (average perceived accuracy rating above the midpoint of the 4-point accuracy scale; 30% of the headlines from fake and hyperpartisan sources). All 120 headlines were reviewed by two coders (for the fact-checking codebook we developed and used, see <https://osf.io/m74v2/>). Of these 120 headlines, 65 were found to be entirely true, 34 were found to contain some falsehood, and no evidence could be found either way for the remaining 21. In other words, less than a third of plausible headlines published by fake and hyperpartisan sources contained demonstrable falsehoods, and roughly twice as many were demonstrably true. Therefore, emphasizing the sources for these headlines would have decreased belief in more *true* headlines than false headlines. That is, emphasizing the sources would have reduced participants' ability to discern truth from falsehood.

Methods

In Study 1, we recruited a target of 600 participants on July 17, 2017 using MTurk. In total, 637 individuals began the study. After excluding those who did not complete the study, responded randomly, searched headlines online during the study, or completed the study twice, the final sample was $N=562$ ($M_{\text{age}} = 35$, $SD_{\text{age}} = 11$, 53.0% female). Although MTurk workers are not nationally representative, they are a reliable resource for research on political ideology (Coppock, 2018; Krupnikov & Levine, 2014; Mullinix, Leeper, Druckman, & Freese, 2015). Furthermore, it is unclear that a nationally representative survey would be more relevant than MTurk, given the target population for Study 1 was people who read and share fake news online.

Participants were randomly assigned to one of three conditions: 1) a no-source condition where no information was provided about the headline's source; 2) a Facebook-baseline condition where the publisher website domain was displayed in small, grey text; and 3) a logo-banner condition where a banner showing the publisher's logo was appended to the bottom of the Facebook post (Figure 1). Publishers' logos were screenshotted from their websites. In all three conditions, participants saw the same set of 12 false headlines and 12 true headlines (taken from Study 1 of Pennycook et al, in press). The headline set contained an equal mix of pro-Republican/anti-Democratic headlines and pro-Democrat/anti-Republican headlines, matched on average pre-tested intensity of partisan slant (see Pennycook et al., in press, for details). Headlines were presented in random order.

As in previous research on fake news (e.g. Bronstein et al., 2019; Pennycook et al., in press; Pennycook & Rand, 2019b,c), our primary dependent measures were an accuracy rating (4-point scale from "Not at all accurate" to "Very accurate") and a sharing intention judgment (3-point scale of "no", "maybe", "yes"). As 80.3% of sharing responses were "no", we dichotomized the sharing variable (0="no", 1="maybe" or "yes"). We deviate from our pre-registered analysis plan of averaging ratings within-subject, because such an approach can inflate the false-positive rate (Judd, Westfall, & Kenny, 2012). Instead, we analyze the results using linear regression with one observation per rating and robust standard errors clustered on subject and headline. (Using logistic regression for the binary sharing decisions does not qualitatively change the results.) We include dummies for the no-source condition and the logo-banner condition, as well as headline veracity (-0.5=false, 0.5=true) and interaction terms. We performed a joint significance

test over the no-source X veracity coefficient and the banner X veracity coefficient to test for an overall interaction between condition and veracity.

Study 2 was identical to Study 1, except (i) the no-source condition was omitted and (ii) the sample size was increased substantially. We recruited a target of 2,000 participants from MTurk on September 6th and 7th, 2017. In total, 2023 participants began the study. After excluding those who did not complete the study, completed the study more than once, or participated in Study 1, the final sample size was $N=1,845$ ($M_{\text{age}} = 35$, $SD_{\text{age}} = 11$, 59% female). The analysis approach was the same as Study 1, but without the no-source dummy and interaction.

For Study 3, we recruited 250 individuals from MTurk ($M_{\text{age}} = 35$, $SD_{\text{age}} = 11$, 36.2% female). Participants were presented with the sources that produced the headlines used in Studies 1 and 2 and asked to make three judgments for each: (i) whether they were familiar or unfamiliar with the source, (ii) whether they trusted or distrusted the source, and (iii) whether they thought the source leaned toward the Democratic or Republican party (or neither). We then calculated the average trust for each source and paired it with the average accuracy rating its headline in the no-source condition of Study 1. Finally, we calculated the correlation across the set of 24 headline-source pairs.

Study 4 consisted of two separate surveys. 1,259 U.S. adults were recruited via MTurk between September 14 and 17, 2018. 251 completed the first survey, and 1,008 completed the second survey (demographics were not collected). (Targets were set at 250 and 1,000, as indicated in our pre-registration.) In the first survey, respondents were presented with a random subset of 30 sources selected from a larger list of 60. This list consisted of 20 “fake-news” sources, 20 hyperpartisan news sources and 20 mainstream news sources and is available in our Supplementary Materials. Respondents indicated their trust in each source using a five-point Likert scale ranging from “Not at all” to “Entirely”.

In the second survey, respondents were presented with a random subset of 30 headlines, which were selected from a larger list of 600 headlines. This list of 600 headlines consisted of the 10 most-popular headlines from each of the 60 sources in the first survey that (i) made claims of fact or opinion and (ii) were published between August 2017 and August 2018. Popularity was established using BuzzSumo and CrowdTangle data. Respondents indicated how accurate they believed each headline to be using a four-point scale ranging from “Not at all accurate” to “Very accurate”.

In Study 5, 1,261 U.S. adults were recruited via Lucid between November 12 and 14, 2018 (253 in the first survey, 1,008 in the second; $M_{\text{age}} = 46$, $SD_{\text{age}} = 30$, 52% female). Lucid is a sample aggregator that uses quota matching to deliver a sample that is nationally representative in age, gender, ethnicity, and geographic region. The experimental design was identical to that of Study 4.

In Study 6, 2,007 U.S. adults were recruited via MTurk on October 23 and 24, 2019 ($M_{\text{age}} = 37$, $SD_{\text{age}} = 11$, 45% female) and asked to rate the accuracy of 30 headlines, which were randomly selected from a subset of 252 headlines from Studies 4 and 5. Obviously outdated headlines were excluded, as were all headlines from Channel 24 News (a fake-news site that succeeded in tricking many participants into believing it was trustworthy). Of the remaining headlines, 70 were rated as accurate in Study 4 (here defined as an average accuracy rating greater than or equal to 2.5 on our 4-point accuracy scale) and were published by a fake-news or hyperpartisan source. Another 144 were likewise published by a fake-news or hyperpartisan source, but were rated as inaccurate in Study 4. We randomly selected 77 of these headlines to include in Study 6. Finally, 105 headlines were published by mainstream sources. Participants were randomly assigned to one of two conditions: 1) a control condition where the text of the headline was presented in isolation and 2) a source condition where the publisher’s domain was indicated in boldface text. Respondents indicated how accurate they believed each headline to be using a four-point scale ranging from “Not at all accurate” to “Very accurate”.

In our final analysis, the 120 headlines from fake and hyperpartisan news sources that received over a 2.5 on our 4-point accuracy scale in both Studies 4 and 5 were reviewed by two raters and categorized as true, false or undetermined according to a shared codebook (<https://osf.io/m74v2/>). Headlines were said

to be true (false) if evidence from credible fact-checkers, credible news organizations or primary sources could confirm (refute) the claims of fact in the headline. The credibility of a news organization or fact-checker was determined using NewsGuard's classifications (NewsGuard, 2019). Disagreements between raters were resolved via discussion. Two headlines could not be agreed upon by raters and were labeled as undetermined.

Bibliography

- Austin, E. W., & Dong, Q. (1994). Source v. Content Effects on Judgments of News Believability. *Journalism Quarterly*, 71(4), 973–983. <https://doi.org/10.1177/107769909407100420>
- Baum, M. A., & Groeling, T. (2009). Shot by the Messenger: Partisan Cues and Public Opinion Regarding National Security and War. *Political Behavior*, 31(2), 157–186. <https://doi.org/10.1007/s11109-008-9074-9>
- Berinsky, A. J. (2017). Rumors and Health Care Reform: Experiments in Political Misinformation. *British Journal of Political Science*, 47(2), 241–262. <https://doi.org/10.1017/S0007123415000186>
- Bronstein, M. V., Pennycook, G., Bear, A., Rand, D. G., & Cannon, T. D. (2019). Belief in Fake News is Associated with Delusionality, Dogmatism, Religious Fundamentalism, and Reduced Analytic Thinking. *Journal of Applied Research in Memory and Cognition*, 8(1), 108–117. <https://doi.org/10.1016/j.jarmac.2018.09.005>
- Coppock, A. (2018). Generalizing from Survey Experiments Conducted on Mechanical Turk: A Replication Approach. *Political Science Research and Methods*, 1–16. <https://doi.org/10.1017/psrm.2018.10>
- Coppock, A., & McClellan, O. A. (2019). Validating the demographic, political, psychological, and experimental results obtained from a new source of online survey respondents. *Research & Politics*, 6(1), 2053168018822174. <https://doi.org/10.1177/2053168018822174>
- Epstein, Z., Pennycook, G., & Rand, D. G. (2019). Will the crowd game the algorithm? Using layperson judgments to combat misinformation on social media by downranking distrusted sources [Preprint]. <https://doi.org/10.31234/osf.io/z3s5k>
- Geoff Samek. (2018, February 2). Greater transparency for users around news broadcasters. Retrieved April 11, 2019, from Official YouTube Blog website: <https://youtube.googleblog.com/2018/02/greater-transparency-for-users-around.html>
- Horton, J. J., Rand, D. G., & Zeckhauser, R. J. (2011). The online laboratory: Conducting experiments in a real labor market. *Experimental Economics*, 14(3), 399–425. <https://doi.org/10.1007/s10683-011-9273-9>
- Jakesch, M., Koren, M., Evtushenko, A., & Naaman, M. (2018). *The Role of Source, Headline and Expressive Responding in Political News Evaluation* (SSRN Scholarly Paper No. ID 3306403). Retrieved from Social Science Research Network website: <https://papers.ssrn.com/abstract=3306403>
- Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology*, 103(1), 54–69. <https://doi.org/10.1037/a0028347>
- Kim, A., Moravec, P. L., & Dennis, A. R. (2019). Combating Fake News on Social Media with Source Ratings: The Effects of User and Expert Reputation Ratings. *Journal of Management Information Systems*, 36(3), 931–968. <https://doi.org/10.1080/07421222.2019.1628921>
- Knight Foundation. (2018, July 18). An online experimental platform to assess trust in the media. Retrieved April 11, 2019, from Knight Foundation website: <https://knightfoundation.org/reports/an-online-experimental-platform-to-assess-trust-in-the-media>

- Krupnikov, Y., & Levine, A. S. (2014). Cross-Sample Comparisons and External Validity. *Journal of Experimental Political Science*, 1(1), 59–80. <https://doi.org/10.1017/xps.2014.7>
- Landrum, A. R., Lull, R. B., Akin, H., Hasell, A., & Jamieson, K. H. (2017). Processing the papal encyclical through perceptual filters: Pope Francis, identity-protective cognition, and climate change concern. *Cognition*, 166, 1–12. <https://doi.org/10.1016/j.cognition.2017.05.015>
- Lazer, D. M. J., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., ... Zittrain, J. L. (2018). The science of fake news. *Science*, 359(6380), 1094–1096. <https://doi.org/10.1126/science.aao2998>
- Mullinix, K. J., Leeper, T. J., Druckman, J. N., & Freese, J. (2015). The Generalizability of Survey Experiments*. *Journal of Experimental Political Science*, 2(2), 109–138. <https://doi.org/10.1017/XPS.2015.19>
- NewsGuard. (n.d.). Retrieved April 11, 2019, from NewsGuard website: <https://www.newsguardtech.com/>
- Pennycook, G., Bear, A., Collins, E., & Rand, D. G. (In press). The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Stories Increases Perceived Accuracy of Stories Without Warnings. *Management Science*.
- Pennycook, G., & Rand, D. G. (2019a). Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proceedings of the National Academy of Sciences*, 116(7), 2521–2526. <https://doi.org/10.1073/pnas.1806781116>
- Pennycook, G., & Rand, D. G. (2019b). Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition*. <https://doi.org/10.1016/j.cognition.2018.06.011>
- Pennycook, G., & Rand, D. G. (2019c). Who falls for fake news? The roles of bullshit receptivity, overclaiming, familiarity, and analytic thinking. *Journal of Personality*. <https://doi.org/10.1111/jopy.12476>
- Pornpitakpan, C. (2004). The Persuasiveness of Source Credibility: A Critical Review of Five Decades' Evidence. *Journal of Applied Social Psychology*, 34(2), 243–281. <https://doi.org/10.1111/j.1559-1816.2004.tb02547.x>
- Rob Leathern. (2018, May 24). Shining a Light on Ads With Political Content. Retrieved April 11, 2019, from <https://newsroom.fb.com/news/2018/05/ads-with-political-content/>
- Sundar, S. S. (2008). *The MAIN Model : A Heuristic Approach to Understanding Technology Effects on Credibility*.
- Swire Briony, Berinsky Adam J., Lewandowsky Stephan, & Ecker Ullrich K. H. (2017). Processing political misinformation: comprehending the Trump phenomenon. *Royal Society Open Science*, 4(3), 160802. <https://doi.org/10.1098/rsos.160802>
- Taylor Hughes, Jeff Smith, & Alex Leavitt. (2018, April 3). Helping People Better Assess the Stories They See in News Feed with the Context Button | Facebook Newsroom. Retrieved April 11, 2019, from <https://newsroom.fb.com/news/2018/04/news-feed-fyi-more-context/>

Funding

The authors gratefully acknowledge funding from the Ethics and Governance of Artificial Intelligence Initiative of the Miami Foundation, the William and Flora Hewlett Foundation, the John Templeton Foundation, and the Social Sciences and Humanities Research Council of Canada.

Competing interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Ethics

All research protocols employed in the course of this study were approved by the institutional review board of either MIT or Yale University. All participants provided informed consent.

Most participants reported race and gender information using categories defined by the investigators. These categories were important to determine how much each sample was representative of the U.S. adult population.

Copyright

This is an open-access article distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided that the original author and source are properly credited.

Data Availability

All materials needed to replicate this study are available via the Center for Open Science:

<https://osf.io/m74v2/>.