

## SOCIAL SCIENCES

# The social media context interferes with truth discernment

Ziv Epstein<sup>1</sup>, Nathaniel Sirlin<sup>2</sup>, Antonio Arechar<sup>2,3</sup>, Gordon Pennycook<sup>4</sup>, David Rand<sup>2\*</sup>

There is widespread concern about misinformation circulating on social media. In particular, many argue that the context of social media itself may make people susceptible to the influence of false claims. Here, we test that claim by asking whether simply considering sharing news on social media reduces the extent to which people discriminate truth from falsehood when judging accuracy. In a large online experiment examining coronavirus disease 2019 (COVID-19) and political news ( $N = 3157$  Americans), we find support for this possibility. When judging the accuracy of headlines, participants were worse at discerning truth from falsehood if they both evaluated accuracy and indicated their sharing intentions, compared to just evaluating accuracy. These results suggest that people may be particularly vulnerable to believing false claims on social media, given that sharing is a core element of what makes social media “social.”

## INTRODUCTION

In recent years, the propagation of misinformation on social media has been a major focus of attention (1–6). Worries about “fake news,” related to both politics and health [e.g., coronavirus disease 2019 (COVID-19)], have led many to see social media as a threat to modern societies (and not, for example, as a tool for promoting collective intelligence and action). Common to such critiques is the assertion that people are more likely to fall for fake news on social media relative to other sources (7–9).

Central to this assertion is the idea that there are distinct affordances to the design of the social media context that may make people particularly susceptible to the influence of false claims [see (7) for a review]. This idea is also captured by common assertions outside the academic literature about the influence of social media algorithms on human psychology, such as in “The Social Dilemma” documentary where it is argued that the “technology that connects us also controls us” (10). Despite this major focus on the psychological impacts of social media, research investigating the causal effects of social media on truth discernment is rather sparse.

Much of the discussion around this topic has focused on features of the social media platforms themselves that may promote misinformation. For example, it has been argued that social media environments often lack salient cues for the epistemic quality of content that exist in other contexts (7). Unlike traditional media, which naturally filters content for veracity with professional gatekeeping, social media is participatory and therefore allows mostly anyone to post (mostly) anything (11, 12). Furthermore, people often engage with news on social media in a manner that is fast-paced and distracted (quickly scrolling through a newsfeed that combines news and emotionally evocative non-news content), and prior work shows that time pressure and distraction (13), as well as emotional arousal (14), can reduce the extent to which people discriminate truth from falsehood. Social media also allows fringe groups to reach large and distributed communities, which, in turn, can

create the illusion that the beliefs of these groups are widespread (15), as well as “false consensus” (15, 16).

Here, we consider an alternative, and in some sense more fundamental, way in which people may be particularly vulnerable to believing fake news on social media. We ask whether simply considering whether or not to share news online interferes with the extent to which people discriminate truth from falsehood when judging the accuracy of news. There are a myriad of motivations for sharing news that go beyond just whether it is accurate (17–20). Because of this, simply making decisions about what to share—and thus be induced to consider these additional motivations—could lead people to be less discerning when asked to judge the news’ accuracy.

There are two distinct (but not mutually exclusive) mechanisms by which asking about sharing could interfere with accuracy judgments. First, it may be that the act of choosing whether or not to share a specific piece of content affects how a social media user perceives the accuracy of that piece of content. In particular, choosing to share a piece of content may make you believe it more (or at least report that you believe it more). Reciprocally, choosing not to share a piece of content might make you believe it less (or at least report that you believe it less). Thus, this first account implies that people prefer to maintain consistency between sharing choices and accuracy judgments.

An alternative mechanism is that thinking about whether to share content is generally distracting from the concept of accuracy by, for example, invoking a social media mindset. That is, simply asking about sharing may cause users to become more noisy when assessing accuracy because thinking about sharing causes people to be distracted by all of the non-accuracy-related motivations and factors that are salient for the sharing choice (17–20). This second account therefore is built on the concept of spillover effects and implies that social media may have an effect because it involves a unique decision-making environment that (at least) fails to prioritize accuracy and truth. Critically, this spillover is not specific to considering whether to share the particular piece of content one is evaluating—rather, even thinking about whether to share the previous piece of content can spill over to influence accuracy judgments of the next piece of content. This is because, for example,

Copyright © 2023 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution NonCommercial License 4.0 (CC BY-NC).

<sup>1</sup>MIT Media Lab, Cambridge, MA, USA. <sup>2</sup>Sloan School of Management, MIT, Cambridge, MA, USA. <sup>3</sup>Center for Research and Teaching in Economics (CIDE), Mexico City, Mexico. <sup>4</sup>University of Regina, Regina, SK, Canada.

\*Corresponding author. Email: drand@mit.edu

considerations relevant to sharing are still top of mind when considering the accuracy of the next piece of content you see.

To test whether deciding what to share interferes with users' accuracy judgments, and to disambiguate between these two potential mechanisms, we conducted an online experiment in which participants were shown a series of true and false headlines. We randomized participants to either (i) simply rate whether each headline is accurate (accuracy-only condition) or (ii) indicate whether they would share each headline on social media and rate its accuracy (and vary the question order, yielding the sharing-accuracy condition and the accuracy-sharing condition; see Fig. 1). This allowed us to test whether the extent to which participants' discriminated true versus false headlines (accuracy discernment) changes when they assess the headlines while also thinking about social media sharing. Last, we also included a condition in which participants are only asked about sharing and not accuracy. This final sharing-only condition allowed us to assess the baseline shareability of each headline absent any accuracy prompt and to test the replicability of past findings that asking about accuracy increases the quality of information people choose to share (21). Our data were collected across two experimental waves: one using text headlines about COVID-19 and a second using headline text + image pairs about politics. The waves also differed in the implementation of the sharing condition: In the first wave, participants were just asked whether they would share each headline, whereas in the second wave, participants were also asked whether they would like and comment on each headline. See Materials and Methods for details.

## RESULTS

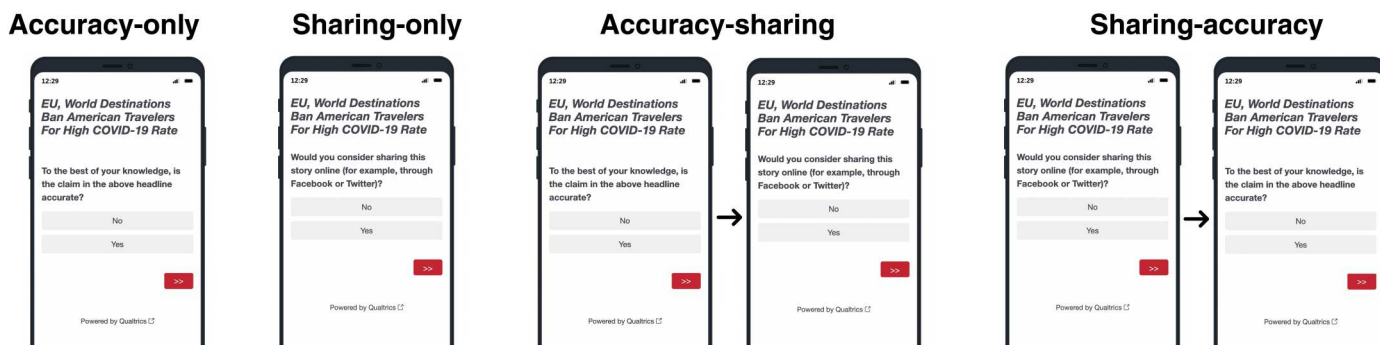
We begin by assessing our key question of interest: How does accuracy discernment (the perceived accuracy of true relative to false headlines) change if participants also make sharing decisions (Fig. 2)? We fit a linear model predicting accuracy rating as a function of headline veracity (0 = false, 1 = true), experimental condition (encoded as per below), wave (z-scored dummy), and all interactions. We encode experimental conditions using two variables: (i) sharing-asked, a dummy variable indicating whether the participant rated sharing and accuracy (sharing-asked, 0 = accuracy-only, 1 = accuracy-sharing or sharing-accuracy), which captures the effect of asking about sharing (collapsing across orders), and (ii) order, a center-coded dummy indicating the order of the two ratings for conditions where both accuracy and sharing were asked ( $-0.5$  = accuracy-sharing,  $0.5$  = sharing-accuracy,  $0$  = accuracy only), which

allows us to test for order effects between the two sharing-asked conditions. Furthermore, because the wave dummy is z-scored, the coefficients on all lower-order interactions can be directly interpreted without needing to run a separate model excluding higher-order interactions with wave.

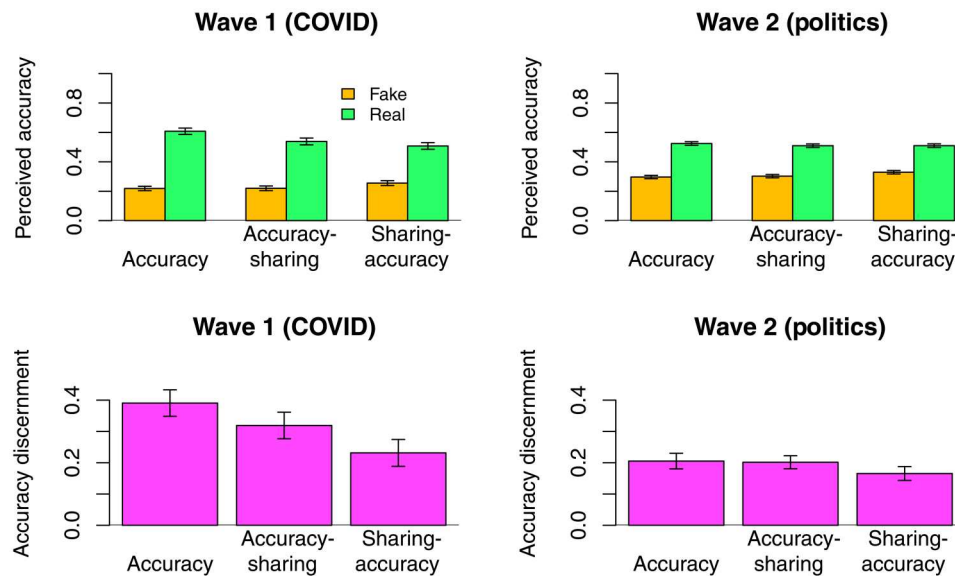
As predicted, the fitted model (the full regression table is shown in table S1) shows a significant decrease in accuracy discernment when sharing is also asked (sharing-asked  $\times$  veracity interaction,  $b = -0.052$ ,  $P < 0.001$ ). Contrary to our preregistered hypothesis, we find that the size of this effect is significantly larger when sharing is asked before accuracy (order  $\times$  veracity interaction,  $b = -0.037$ ,  $P < 0.001$ ). However, we also observe a significant three-way interaction between veracity, wave, and sharing-asked ( $P < 0.001$ ), such that the effect of asking about sharing on accuracy discernment varied across waves. Therefore, we consider the two waves separately.

Considering the first wave (COVID statements with no images or source information), we find a significant main effect of asking about sharing; such accuracy discernment is significantly lower when sharing is also asked (sharing-asked  $\times$  veracity interaction,  $b = -0.103$ ,  $P < 0.001$ ). However, we also find that the size of this effect was significantly different depending on whether accuracy or sharing was asked first (order  $\times$  veracity interaction,  $b = -0.066$ ,  $P = 0.001$ ). Specifically, the decrease in accuracy discernment was substantially larger when sharing was asked before accuracy ( $b = -0.136$ ,  $P < 0.001$ ) compared to when accuracy was asked before sharing ( $b = -0.070$ ,  $P < 0.001$ ). To contextualize these effect sizes, asking about sharing before accuracy led to a 35% decrease in accuracy discernment compared to the accuracy-only baseline, and asking about sharing after accuracy led to a 18% decrease in accuracy discernment relative to the accuracy-only baseline. For the full regression table, see table S2.

Next, we turn our attention to wave 2 (political headlines with image, source information, and liking/commenting as well as sharing in the sharing conditions). Since wave 2 involved political headlines, we also include subject partisanship and political concordance in this model (as per our preregistration). We once again find that, compared to the accuracy-only baseline, accuracy discernment is significantly lower when participants also indicate sharing intentions (sharing-asked  $\times$  veracity interaction,  $b = -0.035$ ,  $P = 0.002$ ). We find that the size of this effect was marginally (but not significantly) different depending on whether accuracy or sharing was asked first (order  $\times$  veracity interaction,  $b = -0.021$ ,  $P = 0.093$ ). That is, the decrease in accuracy discernment is somewhat larger



**Fig. 1. Visualization of different conditions.** For the accuracy-sharing and sharing-accuracy conditions, participants saw the two questions on different pages.



**Fig. 2. Asking about sharing reduces accuracy discernment.** Top: Average perceived accuracy of true (green) and false (orange) news across waves and conditions. Error bars are 95% confidence intervals. Bottom: Subject-level accuracy discernment (perceived accuracy of true minus perceived accuracy of false) across waves and conditions. Error bars are 95% confidence intervals.

when sharing is asked before accuracy ( $b = -0.0455$ ,  $P < 0.001$ ) compared to when accuracy is asked before sharing ( $b = -0.0245$ ,  $P = 0.058$ ). Thus, asking about sharing before accuracy led to a 28% decrease in sharing discernment compared to the accuracy-only baseline (and a nonsignificant 8% decrease for accuracy-then-sharing). We also find little evidence that these effects on accuracy discernment vary significantly based on headline concordance or subject partisanship [although we do find, consistent with past work (22, 23), that preference for the Republican party is associated with lower baseline accuracy discernment,  $b = -0.032$ ,  $P = 0.001$ ]. For the full regression table, see table S3.

Overall, then, our data confirm the prediction that merely considering whether to share news on social media makes people worse at identifying the news' truth—particularly when people consider sharing before judging accuracy.

### Headline-level analysis

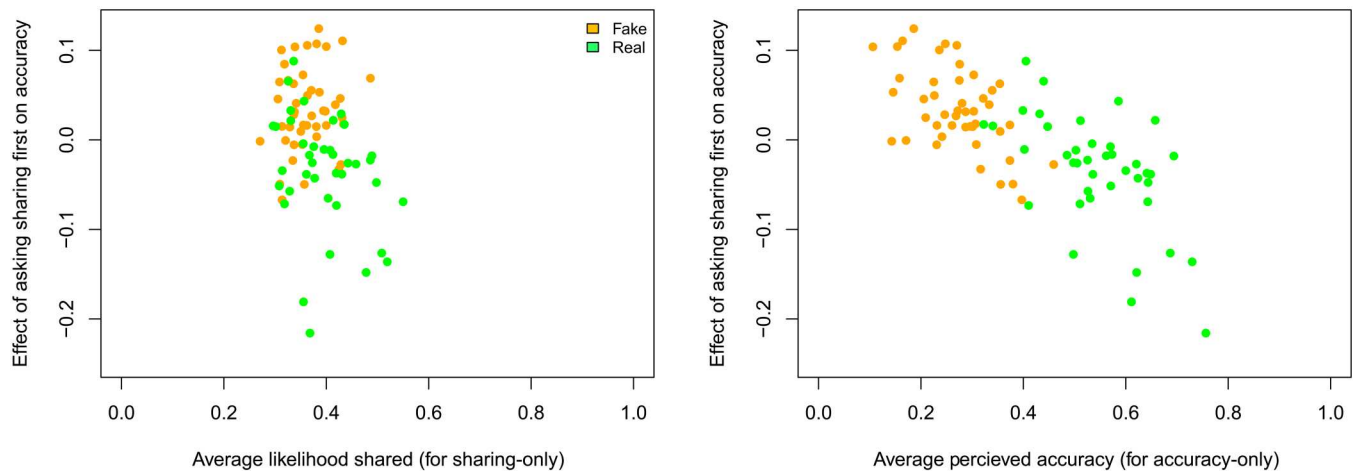
We now aim to shed light on the mechanism underlying this effect. In particular, we seek to differentiate between the consistency-based and spillover-based accounts of why asking about sharing could interfere with accuracy discernment. We do so by evaluating each account's predictions about how the effect of asking about sharing should vary across headlines.

According to the consistency account, sharing a headline makes people think it is more accurate, and/or not sharing a headline makes people think it is less accurate. Thus, the consistency account predicts that, on average, asking about sharing should increase accuracy ratings for headlines people tend to share and/or decrease accuracy ratings for headlines people tend not to share. In other words, the consistency account predicts a positive correlation between (i) how likely a headline is to be shared and (ii) how accuracy judgments of that headline change if sharing is also asked (i.e., perceived accuracy in sharing-accuracy minus perceived accuracy in accuracy-only).

According to the spillover account, asking about sharing should make people less sensitive to accuracy because they are distracted by the many other (non-accuracy-related) factors that are relevant to the choice of what to share. This is posited to be a general mindset effect and thus should occur regardless of how likely the specific headline is to be shared. Instead, the effect of asking about sharing should depend on the baseline perceived accuracy of the headlines—if distraction just adds noise to accuracy ratings, we would expect asking about sharing to decrease the perceived accuracy of headlines that people would otherwise tend to believe and increase belief in headlines that people would otherwise tend to disbelieve. In other words, the spillover account predicts a negative correlation between (i) how accurate a headline is rated in the accuracy-only condition (baseline perceived accuracy) and (ii) how accuracy judgments of that headline change if sharing is also asked (i.e., perceived accuracy in sharing-accuracy minus perceived accuracy in accuracy-only).

To distinguish between these two mechanisms, we conduct two post hoc headline-level analyses (Fig. 3). A regression predicting the effect of asking about sharing on accuracy using independent variables of (i) sharing likelihood in sharing-only, (ii) headline veracity, and (iii) wave finds no significant effect of sharing likelihood in sharing-only ( $P = 0.754$ ; see table S6). This disconfirms the predictions of the consistency account. Conversely, a regression predicting the effect of asking about sharing on accuracy using independent variables of (i) perceived accuracy in accuracy-only, (ii) headline veracity, and (iii) wave finds a significant negative effect of perceived accuracy in accuracy-only ( $P < 0.001$ ; see table S7). This confirms the prediction of the spillover account.

Together, these headline-level analyses support the spillover account over the consistency account. It appears that having users consider whether or not to share headlines generally interferes with their judgments of headline accuracy, making those judgments noisier. This also explains why asking about sharing specifically



**Fig. 3. The size of the asking-about-sharing effect is related to a headline's perceived accuracy, but not its sharability.** Item-level analysis comparing average sharing likelihood in the sharing-only condition (left) and average perceived accuracy in the accuracy-only condition (right) to the difference in perceived accuracy between the sharing-accuracy condition and the accuracy-only condition.

reduces accuracy discernment: Because true headlines tend to be believed more than false headlines, asking about sharing reduces belief in true more than it increases belief in false.

### Moderation by partisanship

Next, we conduct an exploratory analysis assessing whether the effect of asking about sharing on accuracy discernment varies based on participant partisanship. In particular, we inspect the three-way interaction between news item veracity, treatment condition, and preference for the Republican versus Democratic parties (measured using a 1 = strong Democrat to 6 = strong Republican Likert scale). The fitted model (shown in table S8) shows a significant four-way interaction between headline veracity, wave, sharing-asking, and partisanship ( $b = 0.022$ ,  $P = 0.029$ ). Therefore, we consider the two waves separately.

For wave 1 (COVID-19 headlines; regression table shown in table S9), we observe significant moderation of the effect of asking about sharing by political partisanship, as measured by the three-way interaction between veracity, sharing-asked, and partisanship ( $b = -0.036$ ,  $P = 0.047$ ). Specifically, asking about sharing leads to significantly more of a decrease in accuracy discernment for participants who more strongly favor the Republican party over the Democratic party. For wave 2 (political headlines; regression table shown in table S10), conversely, we observe no significant interaction between participant partisanship, headline veracity, and the sharing-asked dummy ( $b = 0.010$ ,  $P = 0.358$ ).

The results are even stronger when we use a binary Democrat versus Republican partisanship measure (by splitting the six-point partisanship Likert scale at the midpoint). Similarly to the continuous partisanship measure, we observe a significant four-way interaction between headline veracity, wave, sharing-asking, and partisanship when using a binary measure of partisanship as well ( $b = 0.026$ ,  $P = 0.006$ ; see table S14), as well as significant moderation of the effect of asking about sharing by political partisanship for wave 1 ( $b = -0.052$ ,  $P = 0.004$ ) and no significant interaction between participant partisanship, headline veracity, and the sharing-asked dummy for wave 2 ( $b = 0.007$ ,  $P = 0.544$ ) (see Fig. 4).

### Sharing intentions

Last, we turn our attention to the effect that asking about accuracy has on sharing intentions. The overall rates of sharing intentions (for both true and false) and discernment for both waves are shown in (Fig. 5). We fit a linear model predicting sharing intention as a function of headline veracity, experimental condition, wave, and all interactions (coded as explained above in the perceived accuracy results). As shown in table S4, we observe no significant interactions with wave and therefore focus on the results pooling across waves.


Consistent with past work (21)—and oppositely from what we observed for the effect of asking about sharing on accuracy—we find that asking about accuracy increases sharing discernment ( $b = 0.021$ ,  $P = 0.017$ ). The size of this effect does not significantly differ based on whether sharing or accuracy was asked first ( $b = -0.001$ ,  $P = 0.93$ )—regardless of order, asking about accuracy translates into a roughly 53% increase in sharing discernment relative to the sharing-only baseline. This means that while putting people in an accuracy mindset increases sharing discernment, putting people in a social media-sharing mindset decreases accuracy discernment. When examining the liking and commenting responses collected in wave 2 as outcomes, the results are not significantly different from the effects observed on sharing: Asking about accuracy increases discernment for sharing, liking, and commenting to similar extents (see table S5).

Exploratory analysis finds no significant moderation of political partisanship on the effect of asking about accuracy on sharing discernment (three-way interaction between veracity, accuracy-asked, and preference for the Republican party,  $b = -0.063$ ,  $P = 0.149$ ) (see Fig. 6 and table S11).

### DISCUSSION

Here, we found that simply asking about social media sharing meaningfully reduces the extent to which people discriminate truth from falsehood when judging headline accuracy. We also find evidence that this occurs because asking about sharing generally distracts people from focusing on accuracy, rather than because



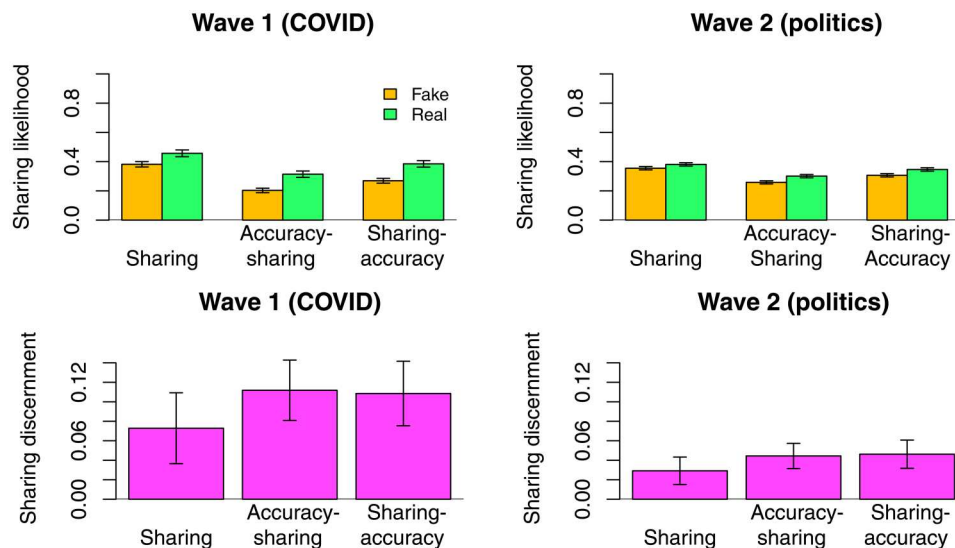


**Fig. 4. Partisanship moderates the asking-about-sharing effect in wave 1 but not wave 2.** Accuracy discernment (perceived accuracy for true headlines minus perceived accuracy for false headlines) by condition across waves and participant partisanship. Error bars are 95% confidence intervals.

people are more likely to believe headlines they want to share (and vice versa). This spillover effect suggests that the social media context—and the mindset that it produces—actively interferes with accuracy discernment. It is not just that people forget to pay attention to accuracy when deciding what to share (21); rather, their actual underlying accuracy judgments are worse when they also consider what to share. Thus, in a similar way that asking about accuracy can induce a more accuracy-focused mindset [and therefore is an actionable intervention to increase sharing discernment (21, 24)], the present results indicate that asking about sharing can induce a mindset focused on social motivations (with harmful, rather than ameliorative, consequences).

Our results on accuracy discernment have important implications for the design of social media platforms. Many platforms center sharing as a principal way users interact with content. Our results suggest that this core design feature may interfere with users' determining what is accurate. Future work should explore alternative design patterns that actively promote truth discernment rather than exclusively privileging engagement (25). This could, for example, be achieved in part by nudging users to consider the

concept of accuracy while scrolling through their newsfeed (21, 24, 26) or redesigning how social cues are displayed (20). A related concern is the context collapse of social media (27) whereby many audiences and types of content are flattened into a single context, which could be mitigated by organizing content and audiences thematically to delineate spaces where accuracy is (e.g., news) versus is not (e.g., family photos) central (28). Alternatively, platforms could emphasize the building of connections between content rather than directly sharing content with an audience. For example, platforms like Are.na (29) and Pinterest achieve this by allowing users to connect or pin content to channels, which does not spread the content but through its relation to user-curated lists. Although we believe that such policies will be net beneficial, there are potential trade-offs to consider. For social media users, these changes may reduce the convenience of the platform by adding friction to interfaces streamlined for sharing and passive consumption; although we believe that this will make the platforms better in the long run, in the short run, it may negatively affect the user experience. For platforms, which are financially motivated to optimize for engagement, there is a different kind of trade-off: Less sharing could



**Fig. 5. Asking about accuracy increases sharing discernment.** Top: Average sharing intentions for true (green) and false (orange) news across waves and conditions. Error bars are 95% confidence intervals. Bottom: Sharing discernment (sharing intentions for true minus sharing intentions for false) across waves and conditions. Error bars are 95% confidence intervals.

lead to higher-quality information, but at the cost of a platform's profits, which may create challenges for getting such policies implemented (30). It is also important to note that our observations have implications beyond just social media: Encouraging sharing is becoming increasingly common across a wide range of online contexts (e.g., news sites, philanthropic organizations, and government agencies), and the negative impacts on accuracy discernment that we observe here may occur in these other contexts as well.

Our results on sharing discernment also have important implications for the deployment of accuracy nudges on social media platforms. While the accuracy nudge has been successfully validated numerous times (21, 24, 26, 31), it remains unclear how forcefully the nudge must be implemented to increase sharing discernment. The lack of order effect that we observed for sharing discernment suggests that it is sufficient to simply prompt an accuracy mindset. In other words, prompting accuracy does not require people to actually write out accuracy for each and every item before sharing (32). This suggests that subtle prompts distributed across a session can effectively reduce the sharing of misinformation on social media. Our replication of the accuracy prompt effect in wave 1, where no source information was provided, also demonstrates that the accuracy prompts are doing more than just making people focus on sources.

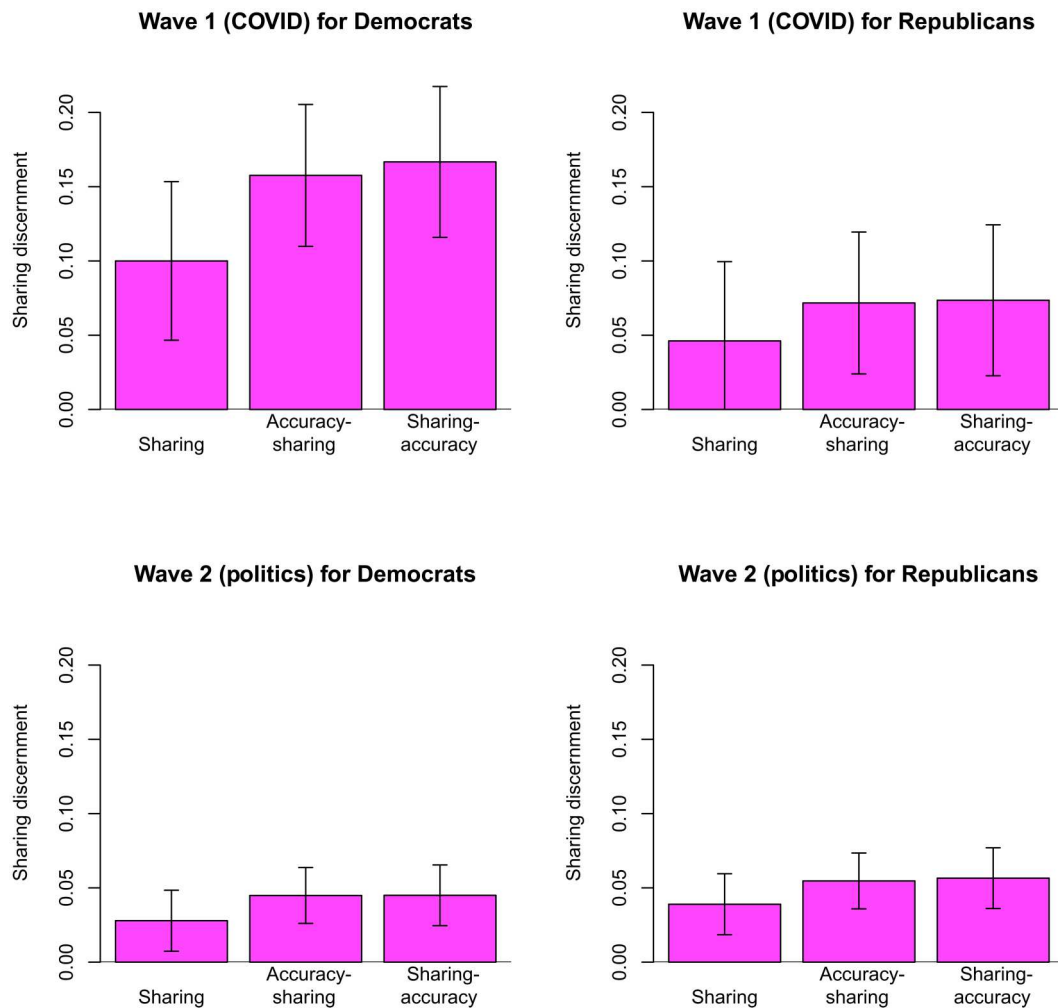
There are several limitations of our study to consider. First, our work uses hypothetical sharing intentions, both as an outcome and as a prompt for the sharing mindset. While previous headline-level analyses have demonstrated that self-reported sharing intentions correlate highly with actual sharing on Twitter and show the same relationships between headline features and sharing rates (33), the actual signifiers of the sharing context are especially important for our findings. Thus, future work should investigate how particular social and design features of a sharing mindset affect truth discernment. Second, our work focuses on the U.S. news ecosystem and recruited U.S. participants. Future work should explore the cross-cultural generalizability of these results.

In sum, we have provided evidence that simply asking about social media sharing reduces truth discernment. This has important implications for social media platforms. Given that eliciting sharing is a (or arguably the) central feature of social media, our results suggest that some level of increased susceptibility to misinformation is inevitable on such platforms. This observation emphasizes the importance of developing interventions that can help users resist falsehoods when online.

## MATERIALS AND METHODS

We conducted two waves of data collection with two similarly structured Qualtrics surveys. Our final combined dataset included a total of  $N = 3157$  participants; mean age = 45.1 years, 59.8% female, 81.3% white. These participants were recruited on Lucid, which uses quota-sampling to approximate the U.S. national distribution on age, gender, ethnicity, and geographic region (34). At the start of the survey, participants were asked "What type of social media accounts do you use (if any)?" and only participants who selected Facebook and/or Twitter were allowed to continue. The studies were exempted by MIT COUHES (protocol 1806400195). Participants gave informed consent before beginning the study.

In both surveys, participants were shown a series of true and false news items. The first wave used a set of 25 headlines (just text) pertaining to the COVID-19 pandemic, 15 false and 10 true. Each participant saw all 25 headlines. We conducted the first study from 29 July 2020 to 8 August 2020 ( $N = 768$ ). The second wave used a set of 60 news "cards" about politics presented in the format of a Facebook post (i.e., including headline, image, and source). These were half true and half false, and participants were shown a random subset of 24 of these headlines. In addition to asking about sharing in the second study, we also asked participants whether they would like or comment on that post ("Would you consider liking or favoriting this story online?" and "Would you consider commenting or replying to this story online?"). We conducted this second study



**Fig. 6. The asking-about-accuracy effect is not moderated by partisanship.** Sharing discernment (sharing intentions for true minus sharing intentions for false) across waves and conditions for Democrats and Republicans. Error bars are 95% confidence intervals.

from 3 October 2020 to 11 October 2020 ( $N = 2389$ ). The full set of headlines, raw data, and analysis code are available at OSF at <https://osf.io/ptvua/>.

As shown in Fig. 1, each participant was randomly assigned to one of four conditions. In the accuracy-only condition, participants were asked for each headline “To the best of your knowledge, is the claim in the above headline accurate?” In the sharing-only condition, they were instead asked “Would you consider sharing this story online (for example, through Facebook or Twitter)?” [as in (24) and (21)]. In the accuracy-sharing condition, participants were first asked the question about accuracy and then, on a separate page, the question about sharing. In the sharing-accuracy condition, participants were first asked the question about sharing and then, on a separate page, the question about accuracy (see Fig. 1). For wave 2, all conditions with the sharing question also contained the questions about liking and commenting. All questions used binary no/yes responses, and we computed both overall (mean) rates of sharing/accuracy and “truth discernment” (i.e., the difference between true and false, with a higher score indicating a greater mean difference). Following the main news item task,

participants then completed a series of demographics and individual difference measures, including political partisanship, a digital literacy battery containing questions about familiarity with internet-related terms and attitudes toward technology (35), a social media literacy question asking how social media platforms decide which news stories to show them (36), and a 10-item procedural news knowledge battery (37). The data from the accuracy-only and sharing-only baseline conditions of our experiment were analyzed in another publication (38), which examined how these individual difference measures predict sharing and accuracy discernment.

For wave 1, the mean accuracy rating (0, 1) was 0.36 with SD 0.48, the mean share rate (0, 1) was 0.33 with SD 0.47, and the mean partisanship (measured using a 1 = strong Democrat to 6 = strong Republican Likert scale) was 3.3 with SD 1.61. For wave 2, the mean accuracy rating was 0.41 with SD 0.49, the mean share rate was 0.32 with SD 0.47, and the mean partisanship was 3.35 with SD 1.74. We completed a preregistration for wave 2 (<https://aspredicted.org/blind.php?x=uy4x3e>). We subsequently realized that our preregistered analysis approach of using multilevel

models and dropping random effects until the model converges is problematic (39). Thus, while we preserved the basic model structure (i.e., which terms were included in the model), we instead used linear regression with two-way clustered errors, the analysis approach used in most of our past work [see, e.g., (21, 40)]. Although we did preregister that we would test for order effects, we note that the order effect was not predicted *ex ante*.

## Supplementary Materials

This PDF file includes:

Sections S1 to S3

Figs. S1 and S2

Tables S1 to S14

## REFERENCES AND NOTES

- S. Vosoughi, D. Roy, S. Aral, The spread of true and false news online. *Science* **359**, 1146–1151 (2018).
- D. M. J. Lazer, M. A. Baum, Y. Benkler, A. J. Berinsky, K. M. Greenhill, F. Menczer, M. J. Metzger, B. Nyhan, G. Pennycook, D. Rothschild, M. Schudson, S. A. Sloman, C. R. Sunstein, E. A. Thorson, D. J. Watts, J. L. Zittrain, The science of fake news. *Science* **359**, 1094–1096 (2018).
- G. Pennycook, D. G. Rand, The psychology of fake news. *Trends Cogn. Sci.* **25**, 388–402 (2021).
- N. Grinberg, K. Joseph, L. Friedland, B. Swire-Thompson, D. Lazer, Fake news on twitter during the 2016 U.S. presidential election. *Science* **363**, 374–378 (2019).
- A. Guess, J. Nagler, J. Tucker, Less than you think: Prevalence and predictors of fake news dissemination on facebook. *Sci. Adv.* **5**, eaau4586 (2019).
- E. Porter, T. J. Wood, The global effectiveness of fact-checking: Evidence from simultaneous experiments in Argentina, Nigeria, South Africa, and the United Kingdom. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2104235118 (2021).
- A. Kozyreva, S. Lewandowsky, R. Hertwig, Citizens versus the internet: Confronting digital challenges with cognitive tools. *Psychol. Sci. Public Interest* **21**, 103–156 (2020).
- S. Bradshaw, P. Howard, The global disinformation order, *Computational Propaganda Research Project* (2019).
- S. Kelly, M. Truong, A. Shahbaz, M. Earp, J. White, Freedom on the net 2017, *Freedom House* (2017).
- J. Orłowski, The social dilemma [documentary film], Los Gatos, CA: Netflix (2020).
- H. Allcott, M. Gentzkow, Social media and fake news in the 2016 election. *J. Econ. Perspect.* **31**, 211–236 (2017).
- M. J. Metzger, A. J. Flanagin, Psychological approaches to credibility assessment online, in *The Handbook of the Psychology of Communication Technology* (John Wiley & Sons, 2015), pp. 445–466.
- B. Bago, D. G. Rand, G. Pennycook, Fake news, fast and slow: Deliberation reduces belief in false (but not true) news headlines. *J. Exp. Psychol. Gen.* **149**, 1608–1613 (2020).
- C. Martel, G. Pennycook, D. G. Rand, Reliance on emotion promotes belief in fake news. *Cogn. Res. Princ. Implic.* **5**, 47 (2020).
- K. Lerman, X. Yan, X.-Z. Wu, The “majority illusion” in social networks. *PLOS ONE* **11**, e0147617 (2016).
- L. Ross, D. Greene, P. House, The “false consensus effect”: An egocentric bias in social perception and attribution processes. *J. Exp. Soc. Psychol.* **13**, 279–301 (1977).
- X. (C.) Chen, G. Pennycook, D. Rand, What makes news sharable on social media? doi:10.31234/osf.io/gzqcd. (2021).
- J. Berger, K. L. Milkman, What makes online content viral? *J. Market. Res.* **49**, 192–205 (2012).
- C. Warren, J. Berger, The influence of humor on sharing, in *Advances in Consumer Research* (Association for Consumer Research, 2011).
- Z. Epstein, H. Lin, G. Pennycook, D. Rand, How many others have shared this? experimentally investigating the effects of social cues on engagement, misinformation, and unpredictability on social media. arXiv:2207.07562 (2022).
- G. Pennycook, Z. Epstein, M. Mosleh, A. A. Arechar, D. Eckles, D. G. Rand, Shifting attention to accuracy can reduce misinformation online. *Nature* **592**, 590–595 (2021).
- G. Pennycook, J. McPhetres, B. Bago, D. G. Rand, Beliefs about COVID-19 in Canada, the United Kingdom, and the United States: A novel test of political polarization and motivated reasoning. *Pers. Soc. Psychol. Bull.* **48**, 750–765 (2022).
- R. K. Garrett, R. M. Bond, Conservatives’ susceptibility to political misperceptions. *Sci. Adv.* **7**, eabf1234 (2021).
- G. Pennycook, D. G. Rand, Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 2521–2526 (2019).
- Z. Epstein, H. Lin, Yourfeed: Towards open science and interoperable systems for social media. arXiv:2207.07478 (2022).
- Z. Epstein, A. J. Berinsky, R. Cole, A. Gully, G. Pennycook, D. G. Rand, Developing an accuracy-prompt toolkit to reduce COVID-19 misinformation online. *Harvard Kennedy School Misinformation Review*, (2021).
- D. Boyd, “Faceted id/entity: Managing representation in a digital world,” thesis, Massachusetts Institute of Technology, Cambridge, MA (2002).
- V. Bhardwaj, C. Martel, D. Rand, Examining accuracy-prompt efficacy in conjunction with visually differentiating news and social online content (PsyArXiv, 2022).
- K. Schwab, “This is what a designer-led social network looks like,” *FastCompany*, 2018.
- K. Hao, “How facebook got addicted to spreading misinformation,” *MIT Technology Review*, 2021.
- A. A. Arechar, J. Allen, A. J. Berinsky, R. Cole, Z. Epstein, K. Garimella, A. Gully, J. G. Lu, R. M. Ross, M. N. Stagnaro, Y. Zhang, G. Pennycook, D. G. Rand, Understanding and reducing online misinformation across 16 countries on six continents (PsyArXiv, 2022).
- L. K. Fazio, Pausing to consider why a headline is true or false can help reduce the sharing of false news. *Harvard Kennedy School Misinformation Review* (2020).
- M. Mosleh, G. Pennycook, D. G. Rand, Self-reported willingness to share political news articles in online surveys correlates with actual sharing on twitter. *PLOS ONE* **15**, e0228882 (2020).
- A. Coppock, O. A. McClellan, Validating the demographic, political, psychological, and experimental results obtained from a new source of online survey respondents. *Res. Polit.* **6**, 2053168018822174 (2019).
- A. Guess, K. Munger, Digital literacy and online political behavior. *Political Sci. Res. Methods* **11**, 110–128 (2023).
- N. Newman, R. Fletcher, A. Kalogeropoulos, D. A. Levy, R. K. Nielsen, Reuters institute digital news report 2017 (2017).
- M. A. Amazeen, E. P. Bucy, Conferring resistance to digital disinformation: The inoculating influence of procedural news knowledge. *J. Broadcast. Electron. Media* **63**, 415–432 (2019).
- N. Sirlin, Z. Epstein, A. A. Arechar, D. G. Rand, Digital literacy is associated with more discerning accuracy judgments but not sharing intentions. *Harvard Kennedy School Misinformation Review*, (2021).
- D. J. Barr, R. Levy, C. Scheepers, H. J. Tily, Random effects structure for confirmatory hypothesis testing: Keep it maximal. *J. Mem. Lang.* **68**, 255–278 (2013).
- G. Pennycook, A. Bear, E. T. Collins, D. G. Rand, 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 (2020).

### Acknowledgments

**Funding:** We gratefully acknowledge funding from the MIT Sloan Latin America Office, the Ethics and Governance of Artificial Intelligence Initiative of the Miami Foundation, the William and Flora Hewlett Foundation, the Reset initiative of Luminate (part of the Omidyar Network), the John Templeton Foundation, the TDF Foundation, the Canadian Institutes of Health Research, the Social Sciences and Humanities Research Council of Canada, the Australian Research Council, and Google. G.P. and D.R. received research support through gifts from Google. G.P. and D.R. received research support through gifts from Facebook. **Author contributions:** Z.E., N.S., and D.R. conceived the research. N.S. and D.R. designed the survey experiments. N.S., A.A., and D.R. conducted the survey experiments. Z.E., N.S., and D.R. analyzed the survey experiments. Z.E., G.P., and D.R. wrote the paper, with input from A.A. and N.S. All authors approved the final manuscript. **Competing interests:** Z.E. has consulted for OpenAI. G.P. is a faculty research fellow at Google. G.P. and D.R. have received research funding from Google, and D.R. has received research funding from Meta. The authors declare that they have no other competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are available online at OSF (<https://osf.io/ptvua/>).

Submitted 27 April 2022

Accepted 18 January 2023

Published 3 March 2023

10.1126/sciadv.abo6169