

Journal of Experimental Psychology: Applied

Quantifying the Effects of Fake News on Behavior: Evidence From a Study of COVID-19 Misinformation

Ciara M. Greene and Gillian Murphy

Online First Publication, June 10, 2021. <http://dx.doi.org/10.1037/xap0000371>

CITATION

Greene, C. M., & Murphy, G. (2021, June 10). Quantifying the Effects of Fake News on Behavior: Evidence From a Study of COVID-19 Misinformation. *Journal of Experimental Psychology: Applied*. Advance online publication. <http://dx.doi.org/10.1037/xap0000371>

Quantifying the Effects of Fake News on Behavior: Evidence From a Study of COVID-19 Misinformation

Ciara M. Greene¹ and Gillian Murphy²
¹ School of Psychology, University College Dublin
² School of Applied Psychology, University College Cork

Previous research has argued that fake news may have grave consequences for health behavior, but surprisingly, no empirical data have been provided to support this assumption. This issue takes on new urgency in the context of the coronavirus pandemic, and the accompanying wave of online misinformation. In this large preregistered study ($N = 3,746$), we investigated the effect of a single exposure to fabricated news stories about COVID-19 on related behavioral intentions. We observed small but measurable effects on some behavioral intentions but not others—for example, participants who read a story about problems with a forthcoming contact-tracing app reported a 5% reduction in willingness to download the app. These data suggest that one-off fake news exposure may have behavioral consequences, though the effects are not large. We also found no effects of providing a general warning about the dangers of online misinformation on response to the fake stories, regardless of the framing of the warning in positive or negative terms. This suggests that generic warnings about online misinformation, such as those used by governments and social media companies, are unlikely to be effective. We conclude with a call for more empirical research on the real-world consequences of fake news.

Public Significance Statement

The coronavirus pandemic has been accompanied by a sharp increase in the quantity of misinformation spreading through social media. This “fake news” is often assumed to have terrible consequences for health behavior, however, no studies to date have directly assessed the effect of specific fake news stories on the targeted behavior. In this study, we created novel fake news stories, suggesting for example, that certain foods might help protect against COVID-19, or that a forthcoming vaccine might not be safe. A single exposure to these fabricated stories resulted in small effects on intentions to engage in some of the behaviors targeted by the stories. For example, we found that participants who saw a fabricated story about privacy concerns with a national contact tracing app were 5% less willing to download the app. We also examined whether providing a general warning about fake news might reduce susceptibility, but found no effects on responses to the fake stories. This suggests that generic warnings such as those often used by governments and social media companies are unlikely to be effective in combating online misinformation. These results are important in our efforts to quantify the real-world effects of fake news and reduce its harms.

Keywords: fake news, misinformation, behavior, warning, COVID-19

Supplemental materials: <https://doi.org/10.1037/xap0000371.supp>

Ciara M. Greene  <https://orcid.org/0000-0002-8833-9046>

Gillian Murphy  <https://orcid.org/0000-0001-8898-139X>

The data and materials associated with this paper may be found at <https://osf.io/mfnb4/>.

We would like to thank Sinéad O’Carroll and Christine Bohan for disseminating this study through The Journal.ie. We would also like to thank Anna Igoe, Jade Duffy, Leigh Huggard and David O’Rourke for assistance with data coding.

Correspondence concerning this article should be addressed to Ciara M. Greene, School of Psychology, University College Dublin, Newman Building, Belfield, D04 V1W8, Ireland. Email: ciara.greene@ucd.ie

Falsehood flies, and the Truth comes limping after it.

Jonathan Swift (1710)

As the COVID-19 pandemic wears on, public health measures instituted across the world have been accompanied by antimask, antilockdown, and antivaccination protests. These movements have spawned much discourse on the causal effects of misinformation, including whether conspiracy theories and “fake news” are the root cause of violent demonstrations (e.g., McDermott, 2021; Rash, 2020; Shackle, 2021). This is a difficult question to answer, as these behaviors are likely the product of many complex, interacting influences, including misinformation. However, this issue does

highlight a surprising gap in the literature—we lack evidence regarding the extent to which exposure to misinformation might affect behavior. The present study, assessing behavior change after a one-off exposure to novel health misinformation in a tightly controlled experimental setting, may help us to bridge this gap and better quantify the threat of real-world misinformation.

Research on fake news has proliferated in recent years, with a particular focus on factors affecting belief in and sharing of fabricated stories (see Lazer et al., 2018 for a review) as well as cognitive processing of misinformation (e.g., Bronstein et al., 2019; Murphy et al., 2019) and interventions aimed at curbing the spread of fake news (Lewandowsky et al., 2012). Much of this research presupposes the grave consequences of fake news exposure, and it is often assumed that online misinformation must be harmful to democracy or public health. There are well-known real-world anecdotes where misinformation has contributed to extreme behaviors, such as the spate of murders in India linked to misinformation circulated via WhatsApp (McLaughlin, 2018), or the “Pizzagate” incident in the U.S., where an armed man entered a pizzeria on a mission to rescue children from a supposed child sex ring that he had read about online (Robb, 2017). Widespread disinformation campaigns can also have demonstrable effects on health behaviors; for example, a drop in the rates of childhood vaccination against measles, mumps, and rubella has been directly tied to debunked misinformation linking the measles, mumps and rubella (MMR) vaccine with autism diagnoses (Leask et al., 2010). Surprisingly, however, there has been little experimental work that has sought to quantify the effect of misinformation exposure on behavior. Given the well-established difficulty of persuading people to change their minds on a given topic (Little, 2018), or of instituting behavioral change in health-related domains (Davis et al., 2015), it is by no means clear that casual fake news exposure *should* be expected to produce significant alterations in health-related behaviors.

We note that the descriptor “fake news” is a catch-all term used to describe inaccurate information, usually spread online via social media and represented as accurate news content. This term incorporates both disinformation (misleading information spread with the intent to cause harm or sow political divisions) and misinformation (misleading information spread unwittingly, without the intention to cause harm; Derakhshan & Wardle, 2017; Wardle & Singerman, 2021). In practice, of course, misinformation and disinformation may overlap: a piece of inaccurate information may at first be intentionally spread to further political ends, before gaining traction or “going viral,” and being shared innocently by members of the public. In the present article, we use the term “fake news” to mean misinformation presented in the form of a news story.

COVID-19 and the Fake News “Infodemic”

The COVID-19 crisis has been accompanied by an increase in the quantity of online misinformation, sparking an “infodemic” (Brennen et al., 2020; Kouzy et al., 2020; World Health Organization, 2020). Anecdotal reports of harms arising from such fake news stories are widespread; for example, it has been argued that exposure to COVID-19 fake news may undermine guidelines on social distancing, or encourage readers to self-medicate with unsanctioned treatments (e.g., Bavel et al., 2020; O’Connor & Murphy, 2020; Tasnim et al., 2020; Wright, 2020; Yoder, 2020). Computer modeling suggests that public response to

disease outbreak might be negatively affected by exposure to misleading health news (Brainard et al., 2020), but no empirical evidence of this has been presented. While some researchers have provided evidence that misinformation may influence beliefs and attitudes (e.g., with respect to climate change; Drummond et al., 2020), we are aware of only one study that has attempted to trace the behavioral effects of fake news in any context. Cantarella et al. (2020) compared Italian municipalities with varying exposure to fake news prior to the 2018 general election, but found no effects on voting patterns. However, as only group-level data were available for this study, inferences about the effects of a given fake news story on individual behavior cannot be drawn. Investigating effects of online misinformation on real-world behavior is undoubtedly difficult, but it is critical that we understand the role of fake news in the current global health crisis. Here, we attempt to address this gap in the literature by providing a straightforward test of the effects of a single exposure to novel COVID-19 misinformation on related behavioral intentions.

There is an abundance of evidence that susceptibility to misinformation can be influenced by a range of factors, including repeated exposure to the information (Pennycook et al., 2018) and whether or not the misinformation aligns with an individual’s preexisting beliefs or social identity (Frenda et al., 2013; Greene et al., 2021; Kahan, 2017; Murphy et al., 2019; O’Connell & Greene, 2017; Van Bavel & Pereira, 2018). Thus, the cognitive and behavioral response of an individual reading a single story about problems with a forthcoming vaccine may be very different to the response of a member of an online “antivax” group who is repeatedly exposed to such misinformation, and who is supported in their suspicions by their peers. In the real world, it can be very difficult to disentangle the cognitive effects of exposure to misinformation from these social pressures. Thus, we restrict our present analysis to novel “fake news” stories, containing misinformation created specifically for this study. While participants will still bring their own social identities and personal beliefs to the experiment, this approach ensures that we can control for any previous exposure to the misinformation, in either a private or social context.

False Memories

Recent research has demonstrated that exposure to fake news can result in the formation of false memories for events that never took place, but were in fact invented by the researchers (Frenda et al., 2013; Greene et al., 2021; Murphy et al., 2019; O’Connell & Greene, 2017; Polage, 2012). For example, Frenda et al. (2013) reported that many participants “remembered” having heard about a fabricated scandal involving President Bush or President Obama after being exposed to a description of the fictional event, accompanied by a doctored photograph. Murphy et al. (2019) found that presenting a fabricated news headline related to the 2018 abortion referendum in Ireland resulted in false memories for political scandals described in the story, while Greene et al. (2021) reported similar findings with regard to fabricated news headlines about Brexit. False memories were more common in each of these studies if the content of the fabricated news stories aligned with participants’ existing political beliefs. In many cases, participants provided detailed descriptions of how they had felt when these scandals were reported in the media, despite the fact that they never actually happened.

False memories can have significant behavioral implications; for example, implanting a false memory of becoming ill as a result of eating a particular food as a child reduces the quantity of that food the participant will choose to eat in the future (Bernstein et al., 2011). Forming a false memory for a fake news story may therefore influence behavioral decisions related to that story. In the context of COVID-19, the consequences of these behavioral choices could be severe, but again, no empirical data have been provided. Thus, we also investigate whether experiencing a false memory for a fake news story has any effects on behavioral intentions over and above mere exposure to the story.

Misinformation Warnings and Corrections

Research on solutions to fake news have mostly focused on targeted fact-checking, which aims to debunk the contents of specific stories. Because misinformation can be “sticky” and difficult to debunk (Lewandowsky et al., 2012; Ross et al., 2018), and fact-checking tags may have the unintended effect of increasing belief in stories that are unlabeled (Pennycook, Bear, et al., 2020), it may be better to prevent acceptance of misinformation in the first place rather than trying to remove it after the fact (Bolsen & Druckman, 2015; Lewandowsky et al., 2020). Some research has experimented with “prebunking,” where participants are warned in advance that the information they are about to see is false (Brashier et al., 2021; Drummond et al., 2020). However, in a situation where misinformation is spreading more rapidly than usual, debunking and fact-checking may not be able to keep up (Bavel et al., 2020), and less targeted approaches may be required. Attempts to protect news consumers from future encounters with fake news have ranged from a simple accuracy nudge (Pennycook, McPhetres, et al., 2020) to gamified “inoculation” interventions (Basol et al., 2020; Roozenbeek & van der Linden, 2019). In an effort to reach wider audiences, some have advocated for simple warnings about misinformation; for example, the U.K. government relaunched their campaign “Don’t Feed the Beast,” to curb the spread of COVID-19 fake news (U.K. Government, 2020), and a radio advertisement in Ireland encouraged the public to “Stop, think, check. Be media smart” (Media Literacy Ireland, 2020). These messages are often positively framed and emphasize gains rather than losses, an approach shown to increase the effectiveness in other domains (Detweiler et al., 1999; Gerend & Maner, 2011; Ivaturi et al., 2017; Rothman et al., 2006). Notably, other research has suggested that loss-framed health information messages, which emphasize the negative consequences of failing to comply with advice, may be more persuasive than gain-framed messages, though the evidence for behavioral effects is weak (see Akl et al., 2011 for a review).

Once again there is little evidence that these general warnings have any effect on belief in fake news stories, or on associated behavioral changes. Clayton et al. (2019) provided their participants with a warning about fake news, pointing out that some stories *in that experiment* might be intentionally misleading. The warning slightly reduced the perceived accuracy of the false headlines, however, participants’ responses may have been influenced by the knowledge that they were participating in a fake-news study. Real-world warnings deployed during the COVID-19 pandemic have been more generic, encouraging critical news consumption in the absence of a specific target. There is little evidence to date about the effectiveness of such pre-exposure warnings.

The Present Study

In the present study, we attempt to quantify the effects of a single exposure to misinformation presented as news headlines on associated behavioral intentions. We focus on fake news stories that might be expected to increase engagement in unhelpful behaviors (e.g., eating more of a certain food) or to reduce engagement in advisable health behaviors (e.g., vaccination). In order to avoid confounds relating to participants’ previous exposure to extant fake news, we presented participants with novel misinformation that had not previously been reported in the media. Before encountering the misinformation, participants were subtly exposed to a generic warning about fake news. We manipulated the framing of the warning in positive or negative terms to compare the effectiveness of gain-framed and loss-framed messaging about misinformation.

We investigated the following preregistered research questions:

1. How does exposure to novel misinformation about COVID-19 or formation of a false memory for this content affect intended health behaviors? Are these effects moderated by exposure to a warning about fake news?
2. How is the perceived truthfulness of a story associated with related behavioral intentions?
3. How do general misinformation warnings affect acceptance of misinformation from fake news stories related to COVID-19?

Method

Preregistration

This study was preregistered at <https://aspredicted.org/de7da.pdf>. The study protocol was approved by the Human Research Ethics Committee (Humanities) at University College Dublin.

Participants

Participants were recruited between May 23 and June 5, 2020 via an article hosted on the Irish news and current affairs website TheJournal.ie (<https://www.thejournal.ie/news-survey-coronavirus-psycho-5105498-May2020/>). Readers of The Journal were informed that psychologists were conducting a study about media coverage of public health messaging, and invited to click a link to an external website to participate. The study was completed by 4,228 participants. In line with our preregistration, 395 people were excluded for failing two attention check questions, and a further 87 were excluded for admitting to seeking help from others or searching for answers on the internet during the study. The final sample included 3,746 participants (1,252 male, 2,487 female, 7 other; M age = 42.29, SD = 12.76, range = 18–101). The majority of participants were well-educated, with 2,395 participants (64%) having earned at least an undergraduate degree. Sixty-nine participants (2.21%) reported having already been infected with COVID-19, and were therefore excluded from analyses relating to behavioral intentions, as outlined in our preregistration. The final sample size provides 95% power to detect small effects ($f < 0.1$) in all analyses described below.

Design

This was a between-subjects design in which participants were randomly assigned to one of four warning conditions: (1) positive misinformation warning, (2) negative misinformation warning, (3) no misinformation warning, and (4) no public health posters.

Materials

All study materials can be found online at <https://osf.io/mfnb4/>.

Public Health Messaging and Misinformation Warning Posters

The posters were designed to mimic the format and style of government-issued public health messages relating to COVID-19 in the Republic of Ireland. Two versions of the misinformation warning poster were developed; the positively framed poster was designed to emphasize the social benefits of careful engagement with suspicious media, while the negatively framed poster emphasized the negative consequences of irresponsible online behavior (see Figure 1). Additional posters presenting standard advice about the importance of social distancing, handwashing, and practicing correct coughing/sneezing etiquette during the COVID-19 crisis may be viewed in Supplemental Materials.

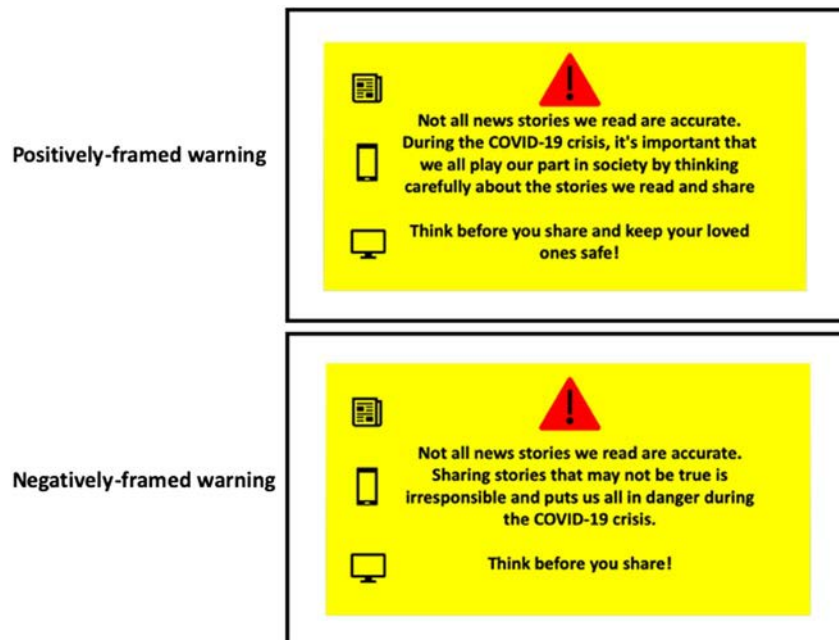
Fabricated Stories

Four novel fake news stories were created for use in this study, each consisting of a short text description accompanied by a nonprobative photograph. The fake stories were designed to be representative of the kind of misinformation typically reported in a

health context (e.g., unproven cures for illnesses and antivax content; Zadrozny, 2019). We explicitly created misinformation that might be expected to result in a change in behavior—for example, eating more of a particular food or declining to participate in public health activities. A wide range of fabricated stories was initially developed, and extensive online research was conducted to determine whether events similar to those described in the fake stories had actually happened or been reported in the media. Any fabricated stories that bore significant similarities to real events were excluded from consideration. We continued to check for any related media reports related to the four selected stories throughout the data collection period, and are confident that the events in question did not occur before or during the experiment. Pilot testing was conducted on 74 participants who did not take part in the final study, and the content of the fabricated stories was amended iteratively in response to pilot participants' estimates of their plausibility (on a scale of 1–9), whether they remembered the events in question and whether they believed the events had happened. The final wording of the four stories used in the experiment was as follows:

1. “New research from Harvard University shows that the chemical in chili peppers that causes the “hot” sensation in your mouth reduces the replication rate of coronaviruses. The researchers are currently investigating whether adding more spicy foods to your diet could help combat COVID-19” (accompanying photograph: a pile of red chili peppers).
2. “A whistleblower report from a leading pharmaceutical company was leaked to the Guardian newspaper in April. The report stated that the coronavirus vaccine being developed by the company causes a high rate of

Figure 1
Positively and Negatively Framed Misinformation Warning Posters



Note. See the online article for the color version of this figure.

complications, but that these concerns were being disregarded in favor of releasing the vaccine quickly” (accompanying photograph: a close-up of a hypodermic needle being inserted into a patient’s arm).

3. “A study conducted in University College London found that those who drank more than three cups of coffee per day were less likely to suffer from severe Coronavirus symptoms. Researchers said they were conducting follow-up studies to better understand the links between caffeine and the immune system” (accompanying photograph: a close-up of a steaming cup of coffee).
4. “The programming team who designed the HSE¹ app to support coronavirus contact-tracing was found to have previously worked with Cambridge Analytica, raising concerns about citizen’s data privacy. The app is designed to monitor people’s movements in order to support the government’s contact-tracing initiative” (accompanying photograph: A close-up of a smartphone with COVID-19 imagery, overlaid on an illustration of the coronavirus organism).

True Stories

Four true stories related to COVID-19 were also shown to participants. These stories were adapted from real news headlines published in the weeks before the start of data collection. As with the false stories, the true stories were each presented alongside a nonprobative photograph. The four stories were as follows:

1. “A new study from Trinity College Dublin revealed that vitamin D is likely to reduce serious coronavirus complications. The researchers urged the government to advise Irish citizens to take daily vitamin D supplements.”
2. “MMA fighter Conor McGregor posted an online video urging the Irish government to enforce a complete lockdown, with the help of the army. “I urge our government to utilize our defence forces” he stated.”
3. “Sinn Féin² President Mary Lou McDonald called off two Sinn Féin rallies in March, after a case of coronavirus was reported at her children’s school.”
4. “As most of Europe is in lockdown, Sweden is pursuing a different strategy against COVID-19. Pubs, restaurants, gyms, and most schools remain open in the Scandinavian state, with the government relying on personal responsibility for compliance rather than strict enforcement. Official guidance states that citizens may socialize, as long as they stay at “arm’s length” from each other.”

Procedure

Participants were informed that the aim of the study was to “investigate reactions to a range of public health messages and news stories relating to the novel coronavirus outbreak”; no reference was made to misinformation or fake news. For ethical reasons, and to reduce the possibility of spreading misinformation during a public health crisis, participants were requested to click a prominent “End Study Now” button rather than simply closing their browser if they wished to halt their participation at any point during the study.

This redirected participants to a full debriefing and an explanation of the truth about the fabricated stories.

A schematic of the experimental procedure may be seen in Figure 2. After consenting to participation, participants provided demographic information (age, gender, and education). They then viewed three study blocks, presented in counterbalanced order. The first block included the warning posters and news stories. Participants in the positively-framed and negatively-framed warning conditions were presented with four public health posters, with the misinformation warning presented in the third position. Participants in the “no misinformation warning” condition viewed the social distancing, coughing, and handwashing posters, but no misinformation posters. Participants were told that these posters were examples of public health messaging that they may have seen during the coronavirus crisis. As any reminder about civic responsibility during the COVID-19 crisis might have the effect of increasing participants’ general level of caution around related news, we also included a fourth condition (“no public health messages”) in which participants viewed no warning posters at all. For each poster, participants were asked to indicate how frequently they had seen a message with similar content since the coronavirus crisis began, using a 5-point scale where 1 = *never* and 5 = *very frequently*. Participants were asked to make this judgment in relation to the content of the poster (e.g., information about the importance of handwashing) and not the appearance of the poster, as the materials used here were created specifically for this study and thus could not have been previously seen by participants.

Immediately after viewing the public health posters, participants were presented with six news stories, including all four true stories and two randomly-selected fake stories. The six stories were presented in random order. After each story, participants were asked, “Do you have a memory of the events described in this story?” and selected a response from the options “I have a clear memory of seeing/hearing about this,” “I have a vague memory of this event occurring,” “I don’t have a memory of this, but it feels familiar,” “I remember this differently” or “I don’t remember this.” Participants were then asked to indicate where they had encountered the story, by selecting all applicable responses from a list (television, newspaper, radio, online news website, social media, word-of-mouth or other source). Participants could also select “I didn’t see/hear about this” or “I don’t remember where I saw or heard about this.” Finally, participants were asked to indicate via an open text box how they felt about the event at the time.

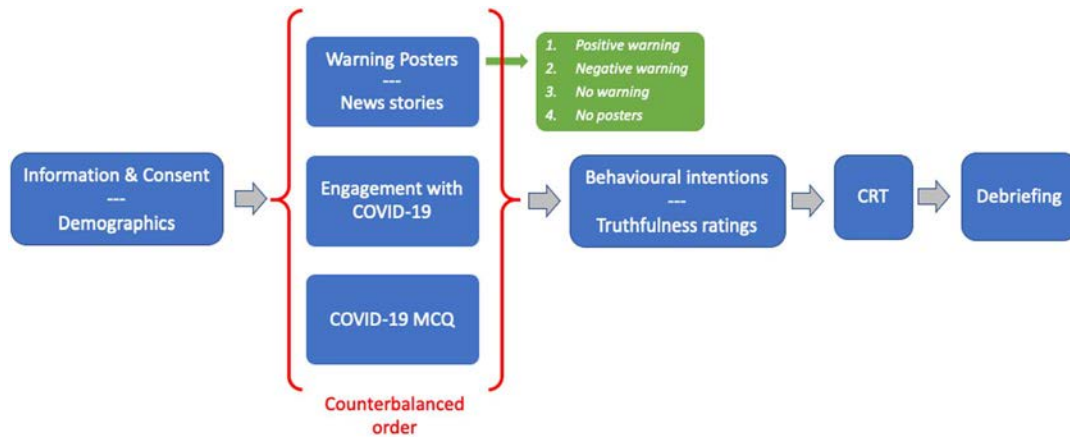
The second block assessed participants’ engagement with news and discussions about the COVID-19 outbreak, while the third block comprised a 10-item multiple-choice quiz assessing knowledge of COVID-19. A final block (presented at the end of the study) invited participants to complete the Cognitive Reflection Test, a measure of analytical reasoning. These variables were collected for a related study assessing individual differences in susceptibility to misinformation, and are reported elsewhere (Greene & Murphy, 2020).

After responding to the news stories and completing the engagement and knowledge questions, participants were asked about their intentions with regard to a list of ten health behaviors over the next several months. Participants were asked to rate each statement on a scale from 1 (*strongly disagree*) to 7 (*strongly agree*). The four critical statements associated with the fake stories read, “I intend to

¹ Health Service Executive, Ireland’s public health service provider.

² Sinn Féin is an Irish political party.

Figure 2
Schematic of Experimental Procedure



Note. See the online article for the color version of this figure.

eat more spicy food,” “I intend to drink more coffee,” “I intend to get a COVID-19 vaccine, once it becomes available,” and “I intend to download the HSE contact-tracing app, once available.” This final statement was accompanied by an explanation that “the app would automatically trace your movements and alert you if you came in contact with a confirmed case of COVID-19.” The remaining statements addressed other health behaviors (e.g., “I intend to get more sleep” and “I intend to get a flu vaccine”). All 10 behavioral intention statements can be found in online Supplemental Materials.

Participants were then informed that “some participants who took part in the study were shown fake news stories (stories concerning events that did not happen, entirely fabricated by the researchers).” A list of the six previously viewed stories (four true and two fake) was then presented, consisting of a brief description of each story and a thumbnail image of the accompanying photograph. Participants were asked to indicate how truthful they considered each story to be on a scale from 0 (*definitely not true*) to 100 (*definitely true*). Each fake story was then presented again, with an explanation that

the story was not true and was fabricated by the researchers. A statement of truth with regard to each fake story was provided (e.g., “there is no evidence that drinking coffee protects against COVID-19”) and was followed by a full debriefing. Note that recent research found no evidence of persistent false memories after taking part in a similar study and being debriefed (Murphy et al., 2020).

Results

Descriptive Statistics

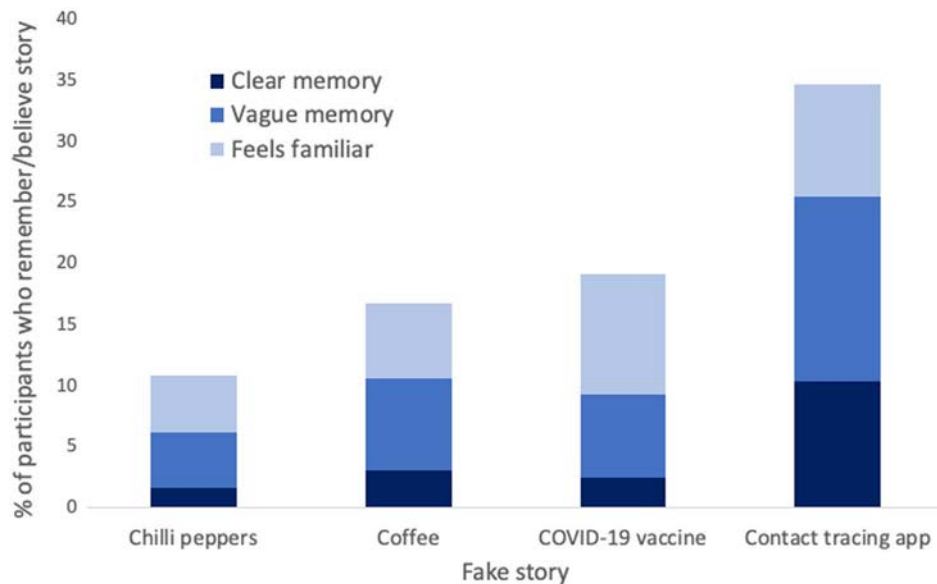
The four critical health behaviors assessed were (a) intention to eat more spicy food, (b) intention to drink more coffee, (c) intention to get a COVID-19 vaccine once available, and (d) intention to download a government-funded contact-tracing app. Each of these behaviors was directly linked to one of the four fake stories described above. Descriptive statistics for the behavioral intentions are listed in Table 1. In total, 845 participants (22.56%) reported a false memory

Table 1
Effect of Warning Condition on Acceptance of Misinformation

Dependent variable	Positive framing <i>M (SD)</i>	Negative framing <i>M (SD)</i>	No misinformation warning <i>M (SD)</i>	No public health posters <i>M (SD)</i>	All conditions <i>M (SD)</i>	Group differences
Behavioral intentions						
Eat more spicy food	2.64 (1.41)	2.80 (1.42)	2.70 (1.46)	2.74 (1.44)	2.72 (1.43)	$F(3, 2974) = 1.59, p = .19, \eta_p^2 = 0.002$
Drink more coffee	2.58 (1.48)	2.64 (1.44)	2.67 (1.50)	2.61 (1.48)	2.62 (1.47)	$F(3, 2971) = 0.46, p = .71, \eta_p^2 = 0.001$
Get COVID-19 vaccine	5.42 (1.76)	5.37 (1.74)	5.29 (1.88)	5.28 (1.84)	5.34 (1.80)	$F(3, 2973) = 1.08, p = .36, \eta_p^2 = 0.001$
Download contact-tracing app	4.75 (2.03)	4.66 (2.04)	4.60 (2.12)	4.74 (1.99)	4.69 (2.05)	$F(3, 2973) = 0.92, p = .43, \eta_p^2 = 0.001$
Truthfulness ratings for fake stories						
Chili peppers	16.12 (21.24)	13.41 (19.07)	16.41 (23.47)	16.66 (22.25)	15.62 (21.54)	$F(3, 1703) = 2.16, p = .09, \eta_p^2 = 0.004$
Coffee	18.36 (24.32)	18.63 (25.05)	17.58 (24.77)	18.20 (23.96)	18.20 (24.50)	$F(3, 1684) = 0.13, p = .94, \eta_p^2 = 0.001$
COVID-19 vaccine	27.94 (29.05)	27.07 (27.29)	25.63 (29.68)	26.59 (26.08)	26.79 (28.08)	$F(3, 1620) = 0.49, p = .69, \eta_p^2 = 0.001$
Contact-tracing app	59.75 (34.83)	59.76 (34.91)	63.43 (34.67)	57.96 (35.41)	60.24 (34.97)	$F(3, 1680) = 1.80, p = .14, \eta_p^2 = 0.003$
False memory count	0.25 (0.48)	0.24 (0.51)	0.25 (0.48)	0.28 (0.52)	0.25 (0.50)	$F(3, 3742) = 0.95, p = .42, \eta_p^2 = 0.001$
Sensitivity (<i>d</i> -prime)	0.01 (1.24)	-0.04 (1.31)	0.06 (1.31)	-0.03 (1.31)	0.00 (1.29)	$F(3, 3742) = 1.07, p = .36, \eta_p^2 = 0.001$
Response bias	-0.01 (1.01)	-0.06 (0.99)	0.04 (0.95)	0.03 (1.05)	0.00 (1.00)	$F(3, 3742) = 1.98, p = .11, \eta_p^2 = 0.002$

Figure 3

Percentage of Participants Who Reported a Clear Memory, Vague Memory, or Sense of Familiarity for Each Fake Story



Note. See the online article for the color version of this figure.

for at least one fabricated story, while 3,665 participants (97.84%) reported a memory for at least one true story. The proportion of participants who remembered each fake story ranged from 6.17% for the chili peppers story to 25.5% for the story about the contact-tracing app. For the true stories, memory rates ranged from 49.41% for the story about mixed-martial arts (MMA) fighter Conor McGregor, to 91.2% for the story about Sweden's response to the coronavirus. Figure 3 depicts the percentage of participants who reported a memory or sense of familiarity for each fake story. In line with our preregistration, a story was considered to be remembered if participants responded either "I have a clear memory of seeing/hearing about this" or "I have a vague memory of this event occurring." Participants reported having seen information similar to the positively framed warning ($M = 2.35$, $SD = 1.30$) slightly more frequently than the negatively framed warning, $M = 2.18$, $SD = 1.21$; $t(1,709) = 2.77$, $p = .006$, $d = 0.13$, and misinformation warnings were seen less frequently than public health advice on coughing, social distancing, or handwashing (all p 's < .001; see Supplemental Materials for details).

Effect of Misinformation on Intended Health Behaviors

Research Question 1 asked how exposure to novel misinformation about COVID-19 or the formation of a false memory for this content affects intended health behaviors, and whether these effects are moderated by exposure to a warning about fake news. As each participant saw two randomly selected fake stories, we can investigate the effects of exposure to a particular story on behavioral intentions. A two-way ANOVA was conducted for each health behavior to investigate the effects of mere exposure to the fake news stories, and determine whether this effect was moderated by a warning about misinformation. To evaluate the effects of false

memories on health behaviors, this was followed by a second analysis among participants who saw each story, examining the effects of remembering versus not remembering the story and warning condition. As this analysis was restricted to participants who had been presented with the story in question, any effects of remembering the event are in addition to the effects of mere exposure to the misinformation. Participants who reported having already been infected with COVID-19 ($N = 69$) were excluded from these analyses, in line with our preregistration. As each of these analyses was conducted four times (once for each behavioral intention), we used a Bonferroni-corrected α of 0.0125 to determine statistical significance.³ The effects of exposure to and memory for the related fake stories on behavioral intentions are depicted in Figure 4, and the effects of warning condition are listed in Table 1.

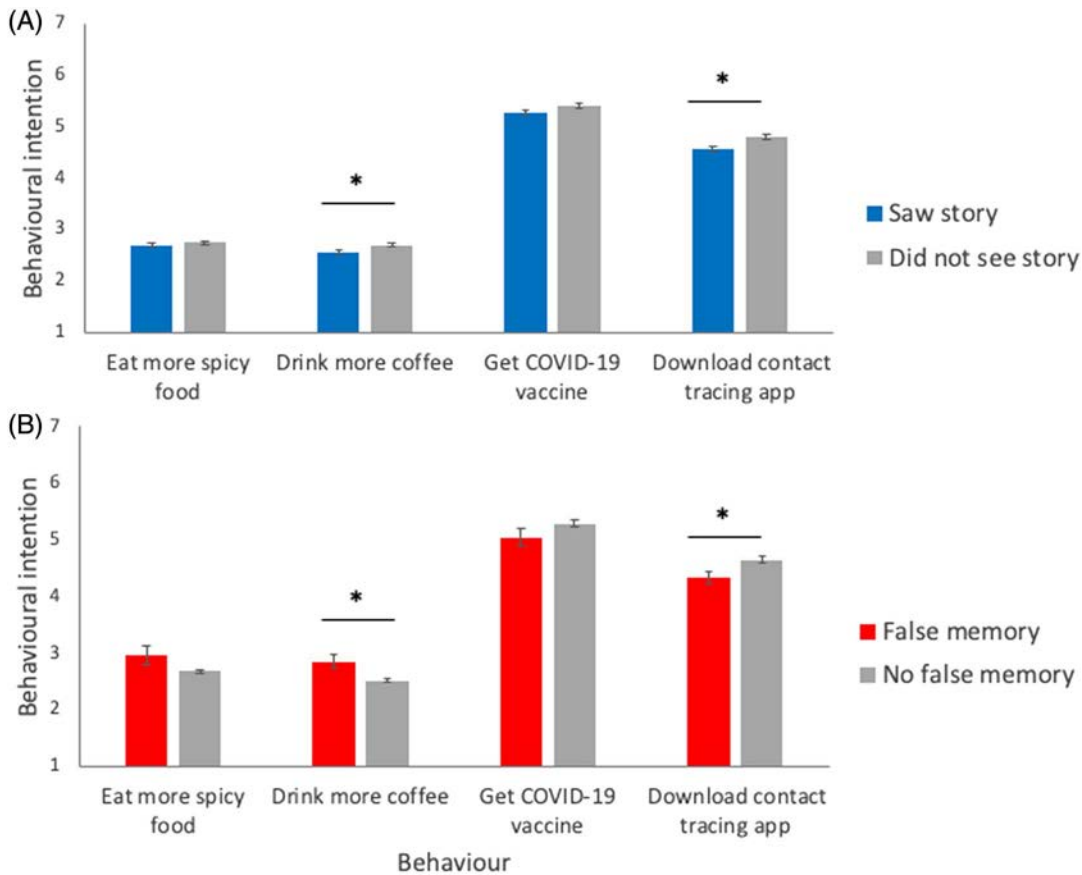
No effect of exposure to the story about chili peppers was observed on intentions to eat more spicy food. Among participants who saw this story, intentions to eat spicy food were slightly higher for those who reported remembering the story ($N = 86$, $M = 2.96$, $SE = 0.15$) compared with those who did not ($N = 1,396$, $M = 2.68$, $SE = 0.04$), but the effect did not reach statistical significance, $F(11, 474) = 3.10$, $p = .08$, $\eta_p^2 = 0.002$. There was no effect of warning condition and no interaction effects in analyses of this behavioral intention (all p 's > .17).

There was a significant effect of exposure to the coffee story on intentions to drink more coffee, though this was in the opposite direction to the expected effect: participants who saw the story about coffee were about 5.5% less likely to agree that they intended to drink

³ This correction was not preregistered; however, on reflection we believe it is appropriate in order to avoid inflated Type I error arising from multiple comparisons across the four critical behavioral intentions.

Figure 4

Effects on Critical Behavioral Intentions of (A) Exposure to and (B) Reporting a False Memory for the Related Fake Story



Note. Error bars represent standard error of the mean. See the online article for the color version of this figure.

* $p < .05$.

more coffee ($M = 2.55$, $SE = 0.04$) than participants who did not see that story, $M = 2.70$, $SE = 0.04$; $F(1, 2967) = 7.45$, $p = .006$, $\eta_p^2 = 0.002$. In contrast, when examining participants who were exposed to the story, a significant effect of memory in the predicted direction was observed, such that participants who reported a memory of the story indicated an 11.6% increase in their intention to drink more coffee ($N = 157$, $M = 2.85$, $SE = 0.12$) compared with those who did not remember the story, $N = 1,315$, $M = 2.52$, $SE = 0.04$, $F(1, 1464) = 7.34$, $p = .007$, $\eta_p^2 = 0.005$. There was no significant effect of warning condition, $F(3, 1464) = 2.33$, $p = .07$, $\eta_p^2 = 0.005$, and no significant interaction effect, $F(3, 1464) = 3.05$, $p = .03$, $\eta_p^2 = 0.006$.

Participants who viewed the vaccine story were approximately 4.5% less willing to get a COVID-19 vaccine ($M = 5.03$, $SE = 0.15$), compared with those who did not view this story ($M = 5.27$, $SE = 0.05$), but the effect did not survive correction for multiple comparisons, $F(1, 2969) = 4.86$, $p = .03$, $\eta_p^2 = 0.002$. There was also no significant difference in behavioral intentions between participants who remembered the events in this story ($N = 142$, $M = 5.05$, $SE = 0.16$) and those who did not, $N = 1,357$, $M = 5.29$, $SE = 0.05$; $F(1, 1491) = 2.34$, $p = .13$, $\eta_p^2 = 0.002$. There were no effects of warning condition and no

interactions between warning condition and exposure to or memory for this story (all $ps > .21$).

Finally, participants who viewed the story about the contact-tracing app reported less willingness to download such an app in future, reporting a 4.8% reduction in intention ($M = 4.57$, $SE = 0.05$) compared with those who did not view this story, $M = 4.80$, $SE = 0.05$; $F(1, 2969) = 9.53$, $p = .002$, $\eta_p^2 = 0.003$. When analysis was restricted to those who were exposed to the story, participants who remembered the events of the story were less willing to download the app ($N = 380$, $M = 4.33$, $SE = 0.11$) than those who did not remember them, $N = 1,117$, $M = 4.65$, $SE = 0.62$; $F(1, 1489) = 6.62$, $p = .01$, $\eta_p^2 = 0.004$, representing a 6.8% reduction in behavioral intention. There was no significant effect of warning condition, $F(3, 1489) = 2.56$, $p = .053$, $\eta_p^2 = 0.005$ and no significant interaction effect, following correction for multiple comparisons, $F(3, 1489) = 3.55$, $p = .014$, $\eta_p^2 = 0.007$.

In summary, exposure to fake news (and subsequent false memory formation) had small effects on some behavioral intentions, but not on others. Viewing a warning about misinformation did not affect behavioral intentions, regardless of framing, and warning condition did not moderate the effects of fake news exposure.

Truthfulness Ratings

After receiving an explicit warning that some of the stories they had just read may have been fake, participants were asked to rate the truthfulness of all six stories to which they were exposed. The false stories were generally rated as less truthful than the true stories, with the exception of the contact-tracing app story which was rated similarly to the four true stories (see Table 1). As might be expected, participants who reported a false memory for a fake story subsequently rated that story as considerably more truthful than those who did not report a memory, chili peppers: remembered $M = 32$, $SD = 33.57$, did not remember $M = 14.57$, $SD = 20.10$, $t(1,705) = 8.11$, $p < .001$, $d = 0.82$; coffee: remembered $M = 41.83$, $SD = 34.75$, did not remember $M = 15.41$, $SD = 21.34$, $t(1,686) = 14.42$, $p < .001$, $d = 1.14$; vaccine: remembered $M = 46.96$, $SD = 34.17$, did not remember $M = 24.83$, $SD = 26.63$, $t(1,622) = 9.26$, $p < .001$, $d = 0.81$; contact-tracing app: remembered $M = 79.15$, $SD = 24.49$, did not remember $M = 53.71$, $SD = 35.67$, $t(1,682) = 13.74$, $p < .001$, $d = 0.76$. There was no effect of warning condition on truthfulness ratings for any of the four stories (see Table 1).

Research Question 2 asked whether the perceived truthfulness of the stories was associated with related behavioral intentions. Small correlations were observed between the truthfulness ratings of the four fake stories, and intention to engage in the related health behavior. As expected, positive relationships between truthfulness and behavioral intentions were observed for the chili peppers story (whereby higher truth ratings were associated with stronger intentions to eat more spicy food, $r(1387) = 0.15$, $p < .001$) and the coffee story (whereby higher truth ratings were associated with stronger intentions to drink more coffee, $r(1381) = 0.10$, $p < .001$). A significant negative relationship was observed for the contact-tracing app, such that higher truth ratings were associated with a reduced intention to download the app, $r(1372) = -0.1$, $p < .001$. A negative relationship was also observed for the vaccine story, but the correlation did not reach statistical significance, $r(1346) = -0.05$, $p = .09$. Thus, participants who believed the fake stories to be more truthful reported stronger intentions to engage in the behavior suggested by the story, but the effects were generally small.

False Memories

On average, participants recalled 0.25 out of 2 fake stories (false memory count; $SD = 0.49$) and 2.65 out of 4 true stories (true memory count; $SD = 1.04$). To further investigate the dynamics of false memory reports, we also computed d -prime, a measure of sensitivity to the difference between true and false stories [calculated as the z score of memories for true stories (“hits”) less the z score of memories for fabricated stories (“false alarms”)], and response bias, a measure of participants’ overall threshold for reporting a memory for any given story (calculated as the z score of “remember” responses). To address Research Question 3 (how do misinformation warnings affect acceptance of misinformation from fake news stories?), one-way Analyses of Variance (ANOVAs) were performed to investigate the effect of warning condition on these memory measures. The details of these analyses may be seen in Table 1; no significant effect of warning condition was observed on false memory count, d -prime, or response bias, despite more than

adequate statistical power. The signal detection measures (d -prime and response bias) are particularly informative, indicating that providing a warning did not increase participants’ ability to discriminate between true and false items, and did not increase their general level of suspicion toward the stories.

To further explore the null effects of warning condition, we conducted a series of equivalence tests (Lakens et al., 2018). These confirmed that mean values on the memory measures, truthfulness ratings, and behavioral intention scales were statistically equivalent across the four warning conditions (see Supplemental Materials for details of these analyses and a full explanation of equivalence testing). We therefore conclude that a general warning about misinformation—whether positively or negatively framed—had no impact on participants’ acceptance of misinformation from fake news.

Discussion

This large, preregistered study aimed to investigate the effects of fake news on health behavior intentions. Exposure to misinformation was associated with small but significant changes to two of the four critical health behaviors assessed. Participants who viewed a story about privacy concerns relating to a contact-tracing app reported being less willing to download the app, while participants who remembered having seen this story before also reported small decreases in intention. A similar effect of exposure to the vaccine story (indicating an approximately 5% reduction in willingness to be vaccinated) was observed, but did not survive correction for multiple comparisons. Participants who reported a false memory for the coffee story reported stronger intentions to drink more coffee in future, though notably the opposite effect was observed among participants who were merely exposed to the story (but did not remember it). No significant effects of seeing or remembering stories about the benefits of eating spicy food were observed; effects were generally in the expected direction, but did not reach statistical significance. Truthfulness ratings were correlated with behavioral intentions; participants who believed stories promoting a particular behavior (e.g., drinking coffee or eating spicy food) tended to report stronger intentions to engage in that behavior. Similarly, participants who believed stories encouraging caution about particular behaviors (e.g., downloading a contact-tracing app or getting a vaccine) were less likely to engage in that behavior in future.

This is the first study to directly evaluate the effects of fake news exposure on behavioral intentions, and represents an important first step in establishing the downstream effects of health misinformation. We report evidence that even a single exposure to health misinformation may “nudge” behavior. It is important to note that we evaluated *intentions* to engage in health behaviors rather than the behaviors themselves. The gap between health intentions and behaviors is well-documented (Rhodes & de Bruijn, 2013; Sheeran, 2002). Intentions are however a critical first step for behavior change: if fake news stories have little effect on targeted behavioral intentions, there is relatively little chance of them impacting actual behavior. The observed effects were relatively small, representing a change of approximately 5%–12% in behavioral intention relative to participants who were not exposed to the fake news story or did not remember the events in question. Nevertheless, it is important to note that small changes in behavior can have significant ramifications, such as when a drop in MMR

vaccinations in the U.K. in the early 2000s led to a reduction in herd immunity, and a corresponding increase in measles cases (Ramsay, 2013). The present study examined relative changes in intention to engage in a behavior, rather than the behavior itself, and it is important to avoid overinterpreting small and inconsistently observed effects. Nevertheless, it is possible that effects of this magnitude could be sufficient to produce measurable harm in the real world. For example, a small reduction in the proportion of the population who choose to download a contact-tracing app could lead to an increase in untraceable transmission of the COVID-19 virus (Hinch et al., 2020).

Changes in behavioral intention were observed in response to some fabricated stories but not others, suggesting that real-world impacts of “fake news” may not be reliably detected. We only assessed four behaviors and it is possible that other behaviors may be easier to influence; perhaps individuals already had sufficiently well-formed opinions about coffee or vaccines that they were resistant to change. Thus, we must acknowledge the possibility that the effects observed in respect of the four fabricated stories (and associated health behaviors) described here may be *sui generis* and may not generalize to other behaviors. While it is fair to note that other topics might be more amenable to the influence of misinformation, a 2019 review of online health misinformation found that the most viral stories were related to cancer, unproven cures, and vaccines (e.g., “Ginger is 10,000× more effective at killing cancer than chemo”; Zadrozny, 2019). Thus we feel our chosen behaviors were reasonably reflective of the online health misinformation that drives concern about public health. A further consideration is the difference between encountering fake news in everyday life compared to seeing these stories in our experiment. Real-world behavioral effects may be expected to arise only following multiple exposures to a story; multiple sources might increase consumers’ faith in a story and thus influence their subsequent behavior. Indeed, just two exposures to a fake news story can increase its perceived truthfulness (Pennycook et al., 2018). Moreover, in real-world settings, fake news is frequently shared among trusted networks; in this setting, it is possible that behavioral effects could be more dramatic and the consequences of following poor health advice consequently more severe. Future research may wish to examine the behavioral effects of multiple exposures to misinformation that appear to come from trusted or untrusted sources.

We also evaluated the effects of positively and negatively-framed warnings about online misinformation. Despite adequate power to detect very small effects, we observed no change in misinformation acceptance or behavioral intention following exposure to a warning, regardless of its framing. This stands in contrast to research suggesting that general warnings about misinformation may be the best means of encouraging critical engagement with suspicious online material (Bolsen & Druckman, 2015; Clayton et al., 2019), and casts doubts on the efforts of governments and social media companies employing similar strategies. This may suggest that, despite the weaknesses identified in previous literature, a targeted strategy focussed on cognitive inoculation or debunking specific stories may be the best available option (Lewandowsky et al., 2020).

It is possible that participants in this study were not influenced by the misinformation warnings because they were already very suspicious about online news content in general, or our study in particular. Participants were not informed that they were taking part in a

misinformation experiment, and thus had no particular reason to be sceptical about the stories presented. Nevertheless, participants reported a lower rate of false memories for fabricated stories than is typical in studies of this kind (Greene et al., 2021; Murphy et al., 2019). We suggest this may reflect the concerted efforts of national and international agencies as well as social media companies to increase awareness of misinformation during the coronavirus crisis (though note that participants reported seeing misinformation warnings much less frequently than other public health advice). As anecdotal support for this conclusion, we note that a number of participants made reference to fake news or misinformation in their responses to both true and false stories, when asked to reflect on how they had felt on first hearing about the events in the story; for example, “Might be a fake news social media story” (51-year-old participant, responding to the chili peppers story), or “It may be true but it’s more likely fake news” (70-year-old participant, responding to the Vitamin D story). Thus, participants’ general level of suspicion when faced with a novel news story may have been increased, leading to a decrease in reported false memories and a reduced impact of the stories on behavioral intention. Notably, participants’ bias toward reporting a memory for any given story was not influenced by warning condition, suggesting that the warnings did not increase this level of suspicion above baseline.

In summary, we report that exposure to misinformation about COVID-19 has small but measurable effects on some health behaviors that are not ameliorated by a general warning. This study represents a first step in quantifying the potential harms of fake news exposure, however, more empirical research is urgently needed to trace the impacts of fake news in the real world. We therefore call on psychological scientists to explicitly measure these effects in order to replace rumor with evidence and allow us to communicate more effectively with the public during times of crisis.

References

- Akl, E. A., Oxman, A. D., Herrin, J., Vist, G. E., Terrenato, I., Sperati, F., Costiniuk, C., Blank, D., & Schünemann, H. (2011). Framing of health information messages. *Cochrane Database of Systematic Reviews*, 2, Article CD006777. <https://doi.org/10.1002/14651858.CD006777.pub2>
- Basol, M., Roozenbeek, J., & van der Linden, S. (2020). Good news about bad news: Gamified inoculation boosts confidence and cognitive immunity against fake news. *Journal of Cognition*, 3(1), Article 2. <https://doi.org/10.5334/joc.91>
- Bavel, J. J. V., Baicker, K., Boggio, P. S., Capraro, V., Cichocka, A., Cikara, M., Crockett, M. J., Crum, A. J., Douglas, K. M., Druckman, J. N., Drury, J., Dube, O., Ellemers, N., Finkel, E. J., Fowler, J. H., Gelfand, M., Han, S., Haslam, S. A., Jetten, J., . . . Willer, R. (2020). Using social and behavioural science to support COVID-19 pandemic response. *Nature Human Behaviour*, 4(5), 460–471. <https://doi.org/10.1038/s41562-020-0884-z>
- Bernstein, D. M., Perna, N. L., & Loftus, E. F. (2011). The false memory diet: False memories alter food preferences. In *Handbook of behavior, food and nutrition* (pp. 1645–1663). Springer. https://doi.org/10.1007/978-0-387-92271-3_107
- Bolsen, T., & Druckman, J. N. (2015). Counteracting the Politicization of Science. *Journal of Communication*, 65(5), 745–769. <https://doi.org/10.1111/jcom.12171>
- Brainard, J., Hunter, P. R., & Hall, I. R. (2020). An agent-based model about the effects of fake news on a norovirus outbreak. *Revue d’Epidemiologie et de Sante Publique*, 68(2), 99–107. <https://doi.org/10.1016/j.respe.2019.12.001>

- Brashier, N. M., Pennycook, G., Berinsky, A. J., & Rand, D. G. (2021). Timing matters when correcting fake news. *Proceedings of the National Academy of Sciences of the United States of America*, 118(5), Article e2020043118. <https://doi.org/10.1073/pnas.2020043118>
- Brennen, J. S., Simon, F. M., Howard, P. N., & Nielsen, R. K. (2020). *Types, sources, and claims of Covid-19 misinformation*. Reuters Institute.
- 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>
- Cantarella, M., Fraccaroli, N., & Volpe, R. (2020). *Does fake news affect voting behaviour?* <https://iris.unimore.it/retrieve/handle/11380/1196186/252258/0146.pdf>
- Clayton, K., Blair, S., Busam, J. A., Forstner, S., Gance, J., Green, G., Kawata, A., Kovvuri, A., Martin, J., & Morgan, E. (2019). Real solutions for fake news? Measuring the effectiveness of general warnings and fact-check tags in reducing belief in false stories on social media. *Political Behavior*, 42(1), 1073–1095. <https://doi.org/10.1007/s11109-019-09533-0>
- Davis, R., Campbell, R., Hildon, Z., Hobbs, L., & Michie, S. (2015). Theories of behaviour and behaviour change across the social and behavioural sciences: A scoping review. *Health Psychology Review*, 9(3), 323–344. <https://doi.org/10.1080/17437199.2014.941722>
- Derakhshan, H., & Wardle, C. (2017). *Information disorder: definitions* [Conference session]. Proceedings of Understanding and addressing the disinformation ecosystem (pp. 5–12). Pennsylvania: Annenberg School for Communication.
- Detweiler, J. B., Bedell, B. T., Salovey, P., Pronin, E., & Rothman, A. J. (1999). Message framing and sunscreen use: Gain-framed messages motivate beach-goers. *Health Psychology*, 18(2), 189–196. <https://doi.org/10.1037/0278-6133.18.2.189>
- Drummond, C., Siegrist, M., & Árvai, J. (2020). Limited effects of exposure to fake news about climate change. *Environmental Research Communications*, 2(8), 081003. <https://doi.org/10.1088/2515-7620/abae77>
- Frenda, S. J., Knowles, E. D., Saletan, W., & Loftus, E. F. (2013). False memories of fabricated political events. *Journal of Experimental Social Psychology*, 49(2), 280–286. <https://doi.org/10.1016/j.jesp.2012.10.013>
- Gerend, M. A., & Maner, J. K. (2011). Fear, anger, fruits, and veggies: Interactive effects of emotion and message framing on health behavior. *Health Psychology*, 30(4), 420–423. <https://doi.org/10.1037/a0021981>
- Greene, C. M., & Murphy, G. (2020). Individual differences in susceptibility to false memories for COVID-19 fake news. *Cognitive Research: Principles and Implications*, 5(1), 63. <https://doi.org/10.1186/s41235-020-00262-1>
- Greene, C. M., Nash, R. A., & Murphy, G. (2021). Misremembering Brexit: Partisan bias and individual predictors of false memories for fake news stories among Brexit voters. *Memory*, 1–18. <https://doi.org/10.1080/09658211.2021.1923754>
- Hinch, R., Probert, W., Nurtay, A., Kendall, M., Wymant, C., Hall, M., Stewart, A., Ferretti, L., Parker, M., Meroueh, A., Mathias, B., Stevenson, S., Montero, D., Warren, J., Mather, N. K., Finkelstein, A., Abeler-Dörner, L., Bonsall, D., & Fraser, C. (2020). *Effective configurations of a digital contact tracing app: A report to NHSX*. https://cdn.theconversation.com/static_files/files/1009/Report_-_Effective_App_Configurations.pdf?1587531217. Retrieved July 23.
- Ivaturi, K., Chua, C., & Janczewski, L. (2017). Impact of information seeking and warning frames on online deception: A quasi-experiment. *Journal of Computer Information Systems*, 57(2), 139–147. <https://doi.org/10.1080/08874417.2016.1183446>
- Kahan, D. M. (2017). *Misconceptions, misinformation, and the logic of identity-protective cognition*. Cultural Cognition Project Working Paper Series No. 164, Yale Law School, Public Law Research Paper No. 605, Yale Law & Economics Research Paper No. 575. <https://doi.org/10.2139/ssrn.2973067>
- Kouzy, R., Abi Jaoude, J., Kraitem, A., El Alam, M. B., Karam, B., Adib, E., Zarka, J., Traboulsi, C., Akl, E. W., & Baddour, K. (2020). Coronavirus goes viral: Quantifying the COVID-19 misinformation epidemic on Twitter. *Cureus*, 12(3), Article e7255. <https://doi.org/10.7759/cureus.7255>
- Lakens, D., Scheel, A. M., & Isager, P. M. (2018). Equivalence Testing for Psychological Research: A Tutorial. *Advances in Methods and Practices in Psychological Science*, 1(2), 259–269. <https://doi.org/10.1177/2515245918770963>
- Lazer, D. M. J., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., & Zittrain, J. L. (2018). The science of fake news. *Science*, 359(6380), 1094–1096. <https://doi.org/10.1126/science.aao2998>
- Leask, J., Booy, R., & McIntyre, P. B. (2010). MMR, Wakefield and The Lancet: What can we learn? *The Medical Journal of Australia*, 193(1), 5–7. <https://doi.org/10.5694/j.1326-5377.2010.tb03730.x>
- Lewandowsky, S., Cook, J., Ecker, U. K. H., Albarracín, D., Amazeen, M. A., Kendeou, P., Lombardi, D., Newman, E. J., Pennycook, G., Porter, E., Rand, D. G., Rapp, D. N., Reifler, J., Roozenbeek, J., Schmid, P., Seifert, C. M., Sinatra, G. M., Swire-Thompson, B., van der Linden, S., ... Zaragoza, M. S. (2020). *The Debunking Handbook*. *Databrary*. <https://doi.org/10.17910/b7.1182>
- Lewandowsky, S., Ecker, U. K. H., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and Its Correction: Continued Influence and Successful Debiasing. *Psychological Science in the Public Interest*, 13(3), 106–131. <https://doi.org/10.1177/1529100612451018>
- Little, A. T. (2018). Fake news, propaganda, and lies can be pervasive even if they aren't persuasive. *Comparative Politics Newsletter*, 28(2), 49–54. https://www.andrewtlittle.com/papers/little_fakenews_cp.pdf
- McDermott, S. (2021, March 3). Adrenochrome: Why was a QAnon conspiracy drug name-checked at Dublin's anti-lockdown protest? *The Journal.ie*. <https://www.thejournal.ie/adrenochrome-dublin-anti-covid-lockdown-protest-qanon-5369654-Mar2021/>
- McLaughlin, T. (2018, December 12). How WhatsApp Fuels Fake News and Violence in India. *Wired*. <https://www.wired.com/story/how-whatsapp-fuels-fake-news-and-violence-in-india/>
- Media Literacy Ireland. (2020). <https://www.bemediasmart.ie/>
- Murphy, G., Loftus, E., Grady, R. H., Levine, L. J., & Greene, C. M. (2020). Fool me twice: How effective is debriefing in false memory studies? *Memory*, 28(7), 938–949. <https://doi.org/10.1080/09658211.2020.1803917>
- Murphy, G., Loftus, E. F., Grady, R. H., Levine, L. J., & Greene, C. M. (2019). False Memories for Fake News During Ireland's Abortion Referendum. *Psychological Science*, 30(10), 1449–1459. <https://doi.org/10.1177/0956797619864887>
- O'Connell, A., & Greene, C. M. (2017). Not strange but not true: Self-reported interest in a topic increases false memory. *Memory*, 25(8), 969–977. <https://doi.org/10.1080/09658211.2016.1237655>
- O'Connor, C., & Murphy, M. (2020). Going viral: Doctors must tackle fake news in the covid-19 pandemic. *BMJ (Clinical Research Ed.)*, 369(369), Article m1587. <https://doi.org/10.1136/bmj.m1587>
- Pennycook, G., Bear, A., Collins, E. T., & Rand, D. G. (2020). The implied truth effect: Attaching warnings to a subset of fake news headlines increases perceived accuracy of headlines without warnings. *Management Science*, 66, 4944–4957. <https://doi.org/10.1287/mnsc.2019.3478>
- Pennycook, G., Cannon, T. D., & Rand, D. G. (2018). Prior exposure increases perceived accuracy of fake news. *Journal of Experimental Psychology: General*, 147(12), 1865–1880. <https://doi.org/10.1037/xge0000465>
- Pennycook, G., McPhetres, J., Zhang, Y., Lu, J. G., & Rand, D. G. (2020). Fighting COVID-19 Misinformation on Social Media: Experimental Evidence for a Scalable Accuracy-Nudge Intervention. *Psychological Science*, 31, 770–780. <https://doi.org/10.1177/0956797620939054>

- Polage, D. C. (2012). Making up History: False Memories of Fake News Stories. *Europe's Journal of Psychology*, 8(2), 245–250. <https://doi.org/10.5964/ejop.v8i2.456>
- Ramsay, M. E. (2013). Measles: The legacy of low vaccine coverage. *Archives of Disease in Childhood*, 98(10), 752–754. <https://doi.org/10.1136/archdischild-2013-304292>
- Rash, W. (2020, June 17). Boogaloo movement tied to murder, violence and disinformation during Protest S. *Forbes*. <https://www.forbes.com/sites/waynerash/2020/06/17/boogaloo-movement-tied-to-murder-violence-and-disinformation-during-protests/?sh=5de4af81552c>
- Rhodes, R. E., & de Bruijn, G. J. (2013). How big is the physical activity intention-behaviour gap? A meta-analysis using the action control framework. *British Journal of Health Psychology*, 18(2), 296–309. <https://doi.org/10.1111/bjhp.12032>
- Robb, A. (2017, November 16). *Anatomy of a fake news scandal*. <https://www.rollingstone.com/feature/anatomy-of-a-fake-news-scandal-125877/>
- Roozenbeek, J., & van der Linden, S. (2019). Fake news game confers psychological resistance against online misinformation. *Palgrave Communications*, 5(1), 1–10. <https://doi.org/10.1057/s41599-019-0279-9>
- Ross, B., Jung, A., Heisel, J., & Stieglitz, S. (2018). *Fake news on social media: The (in) effectiveness of warning messages* [Conference session]. Proceedings of the Thirty Ninth International Conference on Information Systems, San Francisco. <https://aisel.aisnet.org/ficis2018/social/Presentations/16/>
- Rothman, A. J., Bartels, R. D., Wlaschin, J., & Salovey, P. (2006). The strategic use of gain-and loss-framed messages to promote healthy behavior: How theory can inform practice. *Journal of Communication*, 56(Suppl_1), S202–S220. <https://doi.org/10.1111/j.1460-2466.2006.00290.x>
- Shackle, S. (2021, April 8). Among the Covid sceptics: “We are being manipulated, without a shadow of a doubt”. *The Guardian*. <https://www.theguardian.com/news/2021/apr/08/among-covid-sceptics-we-are-being-manipulated-anti-lockdown>
- Sheeran, P. (2002). Intention—behavior relations: A conceptual and empirical review. *European Review of Social Psychology*, 12(1), 1–36. <https://doi.org/10.1080/14792772143000003>
- Tasnim, S., Hossain, M. M., & Mazumder, H. (2020). Impact of rumors or misinformation on COVID-19 in social media. *Journal of Preventive Medicine and Public Health*, 53(3), 171–174. <https://doi.org/10.3961/jpmph.20.094>
- U.K. Government. (2020, March 30). *Government cracks down on spread of false coronavirus information online*. Retrieved July 16, 2020, from <https://www.gov.uk/government/news/government-cracks-down-on-spread-of-false-coronavirus-information-online>
- Van Bavel, J. J., & Pereira, A. (2018). The Partisan Brain: An Identity-Based Model of Political Belief. *Trends in Cognitive Sciences*, 22(3), 213–224. <https://doi.org/10.1016/j.tics.2018.01.004>
- Wardle, C., & Singerman, E. (2021). Too little, too late: Social media companies’ failure to tackle vaccine misinformation poses a real threat. *BMJ (Clinical Research Ed.)*, 372, Article n26. <https://doi.org/10.1136/bmj.n26>
- World Health Organization. (2020, June 29). *Working together to tackle the “infodemic.”* <https://www.euro.who.int/en/health-topics/Health-systems/pages/news/news/2020/6/working-together-to-tackle-the-infodemic>
- Wright, C. L. (2020, April 30). *COVID-19 fake news and its impact on consumers*. <https://www.psychologytoday.com/ie/blog/everyday-media/202004/covid-19-fake-news-and-its-impact-consumers>
- Yoder, K. (2020, July 1). *Fake news is killing us. How can we stop it?* <https://grist.org/climate/fake-news-is-killing-us-how-can-we-stop-it/>
- Zadrozny, B. (2019). *Social media hosted a lot of fake health news this year. Here’s what went most viral*. <https://www.nbcnews.com/news/us-news/social-media-hosted-lot-fake-health-news-year-here-s-n1107466>

Received October 15, 2020

Revision received April 20, 2021

Accepted May 2, 2021 ■