



Perceived social presence reduces fact-checking

Youjung Jun^{a,1}, Rachel Meng^{a,1}, and Gita Venkataramani Johar^{a,1,2}

^aDepartment of Marketing, Columbia Business School, Columbia University, New York, NY 10027

Edited by Susan T. Fiske, Princeton University, Princeton, NJ, and approved April 21, 2017 (received for review January 4, 2017)

Today's media landscape affords people access to richer information than ever before, with many individuals opting to consume content through social channels rather than traditional news sources. Although people frequent social platforms for a variety of reasons, we understand little about the consequences of encountering new information in these contexts, particularly with respect to how content is scrutinized. This research tests how perceiving the presence of others (as on social media platforms) affects the way that individuals evaluate information—in particular, the extent to which they verify ambiguous claims. Eight experiments using incentivized real effort tasks found that people are less likely to fact-check statements when they feel that they are evaluating them in the presence of others compared with when they are evaluating them alone. Inducing vigilance immediately before evaluation increased fact-checking under social settings.

fact-checking | information processing | social influence

Today's media landscape affords people access to richer information than ever before. The diffusion of uncurated digital content into the public sphere, accompanied by the concomitant ease of access to susceptible minds, affords anybody the ability to sway opinions. Exposure to increasingly open communication forums may, in turn, affect the quality of information disseminated. Perhaps nowhere does this concern resonate with greater force than with the emergence of numerous fabricated news stories in the wake of the 2016 US Presidential election, leading some to accuse the internet—and social media in particular—of “distorting our collective grasp on the truth” (1).

The debate over the spread of misinformation acquires broad significance given social media's rise as an important news outlet for many (2, 3), a pattern echoed in our own data (*SI Appendix, News Consumption*). In light of recent coalitions formed between social platform companies and third-party sites to curtail fake news (4), it is useful to ask how individuals scrutinize the content that they encounter on these channels relative to traditional media sources. In this article, we investigate how processing information within a social context affects the extent to which people fact-check ambiguous claims.

Eight experiments suggest that people are less likely to verify statements when they perceive the presence of others, even absent direct social interaction or feedback. The notion of perceived social presence draws from the literature on Social Impact Theory (5, 6) and social facilitation (7), which has examined the influence of noninteractive others—whose “mere presence” may be real, implied, or otherwise imagined—on individual behavior.

Several perspectives on group decision processes inform our prediction of why collective settings might suppress fact-checking. One line of research on social loafing (8) has found that individuals tend to exert less effort in the presence of other coactors, especially if their own input is unidentifiable and dispensable to the group (allowing them to “hide in the crowd”). Although social presence usually leads to loafing behavior for more complex cognitive tasks, such as brainstorming and evaluating proposals, it facilitates simple, effortless tasks, such as motor responses (9, 10). Because the degree of vigilance that fact-checking demands is less likely to be the intuitive response on receiving information, individuals may instead choose to free-ride on others' fact-checking efforts. Such behavior parallels research on diffusion of responsibility and the

bystander effect, whereby people fail to intervene when surrounded by others or simply imagining they are in a group (11).

A second possibility appeals to conversational norms (12–14): In their attempts to make sense of and infer meaning from messages in accordance with social conventions, people often choose by default to take others at their word. This tendency may be exacerbated when others are perceived to be present. Insofar as external fact checks signal skepticism, individuals may be reluctant to express doubt about a speaker's trustworthiness in social situations.

Finally, there could be something inherent—perhaps visceral—about crowds that decreases vigilance, consistent with a “safety in numbers” (i.e., dilution of risk) heuristic observed in animal and human behavior (15, 16). Social Impact Theory has, relatedly, posited that individuals tend to feel less pressure in a social setting because they perceive the impact of an unknown outcome to be proportionally more diffuse as it gets divided among others.

Taken together, the aforementioned forces may, separately or jointly, impede people's willingness to verify information. Note that we distinguish between verification and belief because one's tendency to fact-check a claim can be independent of how true one perceives that claim to be, the latter of which may often be influenced by factors like source credibility (17) and compatibility with existing worldviews (confirmation bias) (18). We did not expect to find systematic social presence effects on individuals' stated beliefs in the statements to which they were exposed; our focus is instead on the verification aspect throughout this paper.* In the experiments that follow, we document a consistent reduction in fact-checking when people perceived the presence of others and seek to discover which explanation(s) can account for this behavior.

Significance

The dissemination of unverified content (e.g., “fake” news) is a societal problem with influence that can acquire tremendous reach when propagated through social networks. This article examines how evaluating information in a social context affects fact-checking behavior. Across eight experiments, people fact-checked less often when they evaluated claims in a collective (e.g., group or social media) compared with an individual setting. Inducing momentary vigilance increased the rate of fact-checking. These findings advance our understanding of whether and when people scrutinize information in social environments. In an era of rapid information diffusion, identifying the conditions under which people are less likely to verify the content that they consume is both conceptually important and practically relevant.

Author contributions: Y.J., R.M., and G.V.J. designed research; Y.J., R.M., and G.V.J. performed research; Y.J. and R.M. analyzed data; and Y.J., R.M., and G.V.J. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹Y.J., R.M., and G.V.J. contributed equally to this work.

²To whom correspondence should be addressed. Email: gvj1@gsb.columbia.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1700175114/-DCSupplemental.

*Social presence did not consistently predict participants' belief in the information (i.e., statements identified as true or false) across experiments (*SI Appendix, Table S2*). Also, actual statement veracity did not systematically affect or interact with social presence to affect the rate of fact-checking (*SI Appendix, Tables S3 and S4*). We thank the reviewers for raising this point.

Methods and Results

Overview. We recruited all respondents online via Amazon Mechanical Turk (MTurk). Participants were unique workers located in the United States with a minimum HIT (Human Intelligence Task) approval rate of 95%. In addition to a fixed payment, respondents could also earn a bonus reward, revealed only after acceptance of the HIT, with specifications that varied across experiments. We excluded data from people (7% of the sample on average) who failed an attention check and indicated that they used outside resources during the focal task. We used the same exclusion criteria throughout, and the results are robust to the inclusion of these participants. *SI Appendix* contains screenshots of stimuli (*SI Appendix, Participant Instructions and Stimuli*) and statements used (*SI Appendix, Statement Selection*) for each experiment.

The rights of our participants were protected throughout the research process, consistent with the regulations set forth by the Institutional Review Board at Columbia University. Informed consent for all studies was obtained by allowing MTurk workers to opt out when they viewed the HIT.

Experiment 1.

Method. In experiment 1 ($n = 175$; $M_{\text{age}} = 36$; 57% female), participants completed a study about (ostensibly) different modes of communication on the internet. Respondents logged onto a simulated news website, where they evaluated 36 statements described as headlines published by a US media organization. We informed people that some headlines were true, whereas others were false, and that their login accounts would be deleted at the end of the study. As validated in a pretest (*SI Appendix, Statement Selection*), these claims were relatively ambiguous, evenly divided in their veracity, and spanned a diverse set of topics.

Participants could identify each statement as true or false; as a third option, they could raise a fact-checking “flag,” which allowed them to learn that statement’s actual veracity at the end. The instructions page delineated the structure of the bonus payment that participants earned depending on their performance. Specifically, we denominated incentives at the statement level, such that people received 1, –1, and 0 points for every correctly identified, incorrectly identified, and flagged statement, respectively (1 point = \$0.05). Hence, a given individual could receive a total bonus amount ranging from \$0 to \$1.80.

For the duration of the evaluation task, some participants saw their username displayed by itself on each screen (alone), whereas others saw the names of 102 other “currently logged on” participants beneath their own (group). After providing responses (true, false, or flag) for 36 statements, participants responded to several measures pertaining to their experience, including how confident they felt (1 = not at all; 10 = extremely) and the extent to which they were influenced by other study respondents (0 = not at all; 10 = extremely).

Results. Participants flagged, or fact-checked, fewer statements in the group compared to the alone condition [$M_{\text{Alone}} = 6.06$ (17% of total statements), $SD = 7.00$; $M_{\text{Group}} = 4.01$ (11%), $SD = 4.92$; $F(1,173) = 5.01$, $P = 0.03$].[†] This pattern was not driven by asymmetric confidence levels because both groups reported feeling equally confident about their answers throughout the task [$M_{\text{Alone}} = 6.16$, $SD = 2.56$; $M_{\text{Group}} = 6.47$, $SD = 2.15$; $F(1,173) = 0.73$, $P = 0.39$]. Also, people did not seem to have much insight into their own behavior, with those in the group condition steeply discounting how much they would be swayed by others’ presence [$M_{\text{Alone}} = 0.78$; $M_{\text{Group}} = 1.07$; $F(1,173) = 1.19$; $P = 0.28$].

[†]Because flagging responses were not normally distributed, we analyzed these responses using log-transformed data in all the experiments. These analyses did not produce qualitatively different results from the ones using the raw data reported herein (for ease of interpretability).

Experiment 2.

Method. The incentive structure in experiment 1 implies that an expected value maximizer would be indifferent between flagging and not flagging (under complete ambiguity, guessing true/false and selecting the flag option both yield an expected value of \$0). To test whether people are more likely to fact-check in the presence of others when doing so is strictly dominant, experiment 2 ($n = 215$; $M_{\text{age}} = 36$; 55% female) imposed a small reward for flagging. Specifically, participants received 1, –1, and 0.25 points for every correctly identified, incorrectly identified, and flagged statement, respectively (1 point = \$0.05). To minimize potential ceiling effects given the dominance of fact-checking, we created a mixture of 37 statements (“news headlines”) that varied in perceived ambiguity, incorporating a subset of 12 statements pretested to be fairly unambiguous.

Results. Although overall flagging rates increased because of the dominance of this strategy, participants again fact-checked fewer statements in others’ presence [$M_{\text{Alone}} = 9.74$ (26%), $SD = 8.09$; $M_{\text{Group}} = 7.29$ (20%), $SD = 6.55$; $F(1,213) = 5.94$, $P = 0.02$] (*SI Appendix, Fig. S1*). This pattern was true irrespective of a claim’s ambiguity. An analysis regressing flagging rates on social presence (0 = alone; 1 = group), clarity indices (a measure of inverse ambiguity) (*SI Appendix, Statement Selection*), and their interaction revealed that a statement’s clarity level did not interact with social presence [$B = -0.03$, $SE = 1.07$, $t(70) = -0.03$, $P = 0.98$]. Although overall fact-checking frequency decreased with statement clarity [i.e., lower ambiguity; $B = -2.24$, $SE = 0.76$, $t(70) = -2.96$, $P = 0.004$], people fact-checked less often in the group (vs. alone) condition across the board.

Experiment 3.

Method. Experiment 3 investigated the diffusion of responsibility account for lowered fact-checking under social presence. Respondents ($n = 165$; $M_{\text{age}} = 36$; 51% female) read 38 statements about real US congressmen/-women as part of a study simulating an online political forum. We added a third condition (group-distinct) that differentiated participants from others ostensibly logged onto the forum. Specifically, we heightened individual distinctiveness by displaying the participant’s name in red text alongside the names of 30 others in black, a method used in research on self-signaling and attention (19).

To ensure the possibility of social loafing, we informed participants in the group and group-distinct conditions that they would be working as a team with the other respondents logged in, such that success on the task would depend on the sum of each individual’s score. If perceived presence blunts fact-checking because individuals believe that they can free-ride on the efforts of others, then this behavior should diminish when they feel distinct, and hence individually responsible, within a collective setting. Participants in the group (vs. alone) conditions entered a lottery for a \$100 Amazon gift certificate if their team (vs. they) scored within the top 10% of respondents. We further imposed a 0.25-point penalty for flagging to better mirror the costs (e.g., search costs) associated with fact-checking in everyday decisions. As with previous experiments, correctly identified and incorrectly identified statements received 1 and –1 point, respectively.

After the evaluation task, participants reported how much they felt like they were part of a group and how responsible they felt about discovering the veracity of the statements. Those in the two group conditions further reported how much they felt like their presence in the forum was noticeable and prominent (all measures scaled from 1 = not at all to 7 = very much).

Results. Participants in both group conditions felt a greater sense of collective membership compared with those in the alone condition [$M_{\text{Alone}} = 2.00$, $SD = 1.34$; $M_{\text{Group}} = 2.98$, $SD = 2.05$; $M_{\text{Group-Distinct}} = 2.64$, $SD = 2.67$; $F(1,162) = 6.25$, $P = 0.01$] but did not differ from each other [$F(1,162) = 0.45$, $P = 0.50$]. Those in the group-distinct condition who saw their name highlighted in red reported feeling more noticeable and prominent in the forum compared with their counterparts in the group condition

$[M_{\text{Group-Distinct}} = 4.03, SD = 2.34; M_{\text{Group}} = 2.51, SD = 1.70; F(1,106) = 14.57, P < 0.001]$. Furthermore, these participants reported equal levels of responsibility as those in the alone condition $[M_{\text{Group-Distinct}} = 5.16, SD = 1.64; M_{\text{Alone}} = 4.53, SD = 2.01; F(1,162) = 1.80, P = 0.54]$ and more than those in the group condition $[M_{\text{Group}} = 3.98, SD = 1.94; F(1,162) = 10.7, P = 0.001]$.

However, these feelings in the group-distinct condition did not increase fact-checking. Flagging rates differed by condition $[F(2,162) = 3.18, P = 0.04]$. Those in the alone condition $[M = 5.02 (13.2\%), SD = 7.72]$ flagged more frequently compared with those in both the group condition $[M = 2.25 (5.9\%), SD = 4.61; F(1,162) = 5.55, P = 0.02]$ and the group-distinct condition $[M = 2.76 (7.3\%), SD = 5.62; F(1,162) = 3.72, P = 0.05]$; the latter two did not differ from each other $[F(1,162) = 0.21, P = 0.65]$. Taken together, the finding that individual distinctiveness increased feelings of responsibility on a collective task but did not increase fact-checking casts doubt on a diffusion of responsibility explanation.

Experiment 4.

Method. So far, we have represented social presence by displaying the names of other respondents presumed to be completing the same task. Experiment 4 ($n = 371; M_{\text{age}} = 36; 56\%$ female) tested whether evaluating information in a social context where the presence of others is cued indirectly (e.g., as when browsing a social media platform) can similarly impede fact-checking.

Participants evaluated the same 36 headlines from experiment 1 (under similar incentive specifications) while viewing their own name by itself (alone) or alongside those of 102 others (group). We introduced, between subjects, a second factor of platform type: those assigned to the traditional news condition completed the evaluation task on the same simulated news website as in previous experiments, whereas those in the social media condition viewed the same headlines designed to appear as Facebook posts from the same focal media organization. Afterward, people evaluated how much they felt like they were in a group during the task (1 = not at all; 7 = very much).

Results. Regardless of whether they saw their name alone or together with others, participants exposed to statements presented as Facebook posts reported similar levels of group belonging $[M_{\text{Alone}} = 2.01, SD = 1.57; M_{\text{Group}} = 2.23, SD = 1.71; F(1,367) = 0.89, P = 0.34]$. However, those in the traditional news condition expressed greater group membership only when they saw the names of others $[M_{\text{Alone}} = 1.75, SD = 1.44; M_{\text{Group}} = 2.30, SD = 1.70; F(1,367) = 5.37, P = 0.02]$.

A 2 (social presence) \times 2 (platform) ANOVA on the number of flagged statements (Fig. 1) revealed a social presence \times platform interaction $[F(1,367) = 3.31, P = 0.07]$. Replicating prior results, people flagged fewer statements when they saw others (vs. only themselves) on the traditional media site $[M_{\text{Alone}} = 2.33 (6.4\%), SD = 3.61; M_{\text{Group}} = 1.27 (3.5\%), SD = 2.36; F(1,367) = 6.38, P = 0.01]$. On the social media platform, however, this difference disappeared $[M_{\text{Alone}} = 1.58 (4.4\%), SD = 2.57; M_{\text{Group}} = 1.62 (4.4\%), SD = 2.74; F(1,367) = 0.01, P = 0.93]$. Those who completed the task alone on the traditional website flagged more statements than those in all other conditions $[F(1,367) = 5.90, P = 0.02]$. These patterns suggest that viewing information on social media—absent direct cues about other individuals—sufficed to reduce fact-checking.

Experiment 5.

Method. As indicated by participants' open-ended inferences,[‡] the task that we have administered thus far led those in the group conditions to believe that they were evaluating the same statements

[‡]In experiment 4, participants described what they thought the other respondents logged onto the website were doing during the task. Two coders independently evaluated 185 free responses from those in the group conditions ($\kappa = 0.83, 95\%$ confidence interval = 0.79, 0.87). A majority (76%) believed that others were performing the same task concurrently, 10% thought they were not doing the same task, and 14% indicated that they were not sure.

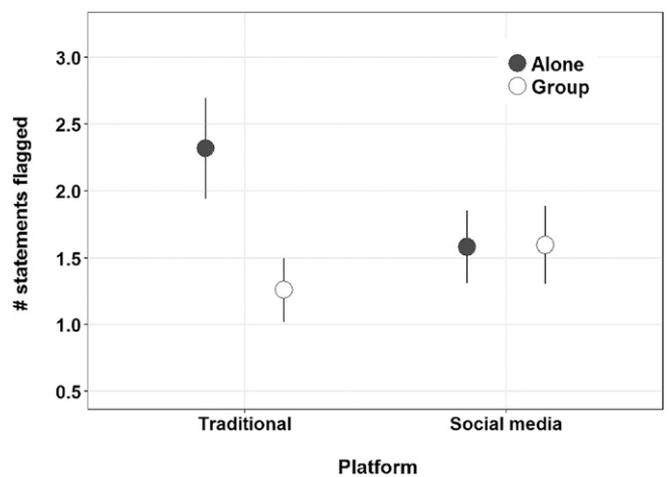


Fig. 1. Number of statements flagged (i.e., amount of fact-checking) by social presence and platform (experiment 4). Error bars denote SEs (\pm SEM).

concurrently with others (i.e., shared attention) (20). Experiment 5 ($n = 308; M_{\text{age}} = 35; 45\%$ female) assessed whether people are more likely to fact-check when they perceive others' presence but realize that they are not coattending to the same stimuli. This scenario echoes situations that individuals frequently face on social media when they observe previous reactions to a post as opposed to responses in real time.

Respondents evaluated 36 headlines (as in experiment 1, with identical incentives) under one of three conditions: alone; group-present, where they saw the names of 102 other "currently online" participants described to be completing the same task simultaneously; and group-past, where they saw 102 other "previously online" participants who had completed the task 1 week ago. Then, everyone reported how much they felt like they were in a group during the task (1 = not at all; 7 = very much).

Results. Flagging rates varied by condition $[F(2,305) = 4.83, P = 0.01]$. Compared with those whose names appeared alone $[M = 8.80 (24\%), SD = 9.32]$, participants who evaluated statements in the company of both current others $[M = 5.47 (15\%), SD = 7.19; F(1,305) = 9.36, P = 0.002]$ and past others $[M = 6.62 (18\%), SD = 6.80; F(1,305) = 3.92, P = 0.04]$ flagged less frequently; these groups did not differ from each other $[F(1,305) = 1.08, P = 0.30]$. Those who saw current or past participants felt equivalent levels of group membership compared with those who only saw their own name $[M_{\text{Group-Present}} = 2.55, SD = 1.93; M_{\text{Group-Past}} = 2.73, SD = 1.91; M_{\text{Alone}} = 1.95, SD = 1.45; F(1,305) = 10.4, P = 0.001]$, suggesting that others need not be simultaneously engaged with the same stimuli to make their presence felt.

Experiment 6.

Method. The news headlines in experiments 1–5, albeit diverse in content, were predominantly neutral in tone. Using more polarizing claims, experiment 6 ($n = 287; M_{\text{age}} = 37; 56\%$ female) tested whether the effect of social presence on fact-checking operates independently from people's motive to trust information that corroborates their own beliefs. Participants read 50 campaign statements made by two US politicians (candidates A and B). These statements were described as posts that the politicians shared in a political forum before an election. Candidate A's statements reflected a more conservative view, whereas candidate B's statements reflected a more liberal one, but their political affiliations were not explicitly mentioned. We counterbalanced the presentation order of candidates and randomly ordered 25 statements that each candidate made. Because people generally do not expect politicians to be very trustworthy (*SI Appendix, Source Credibility*), they

should be less likely to deem fact-checking as a socially inappropriate behavior in this context.

We further varied the number of people logged onto the forum during the evaluation task to examine the role of group size on fact-checking. Respondents completed the task while seeing their own name displayed by itself (alone), alongside those of 30 others (group-small), or alongside those of 102 others (group-large). After evaluating the statements, participants reported their own political affiliation (Republican, Democrat, Independent, or other) as well as how likely they thought it was that the candidates would lie (1 = not at all; 7 = very much). If participants believe that the candidates are likely to lie, they should feel less bound by conversational norms and therefore freer to express skepticism through fact-checking.

Results. Flagging rates varied by condition [$F(2,284) = 3.82, P = 0.02$], with participants fact-checking less under the presence of others, regardless of group size [$F(1,284) = 7.18, P = 0.01$]. Compared with those in the alone condition [$M = 5.82$ (12%), $SD = 8.34$], participants who made their decisions in both a smaller group [$M = 3.78$ (8%), $SD = 6.45$; $F(1,284) = 3.88, P = 0.05$] and a larger one [$M = 3.03$ (6%), $SD = 6.70$; $F(1,284) = 7.07, P = 0.01$] flagged fewer statements; flagging rates did not vary according to group size [$F(1,284) = 0.51, P = 0.48$].

Fig. 2 illustrates the proportions of true, false, and flagged responses across participants by candidate affiliation (liberal vs. conservative) and own party identification (Democrat vs. Republican). To test for the presence of confirmation bias, we coded political alignment such that participants whose party identification matched the views expressed by a candidate were coded as one and zero otherwise. A 2 (alignment) \times 3 (social presence) ANOVA on the responses of 184 participants who identified as either Democrat or Republican⁵ revealed that people were more readily persuaded to believe in claims consistent with their own political worldviews, but this tendency did not depend on whether they evaluated statements alone or in a group.

Specifically, Democrats evaluated a liberal candidate's statements as more true [$M_{Dem} = 55\%$ vs. $M_{Rep} = 43\%$; $F(1,178) = 17.4, P < 0.001$] and less false [$M_{Dem} = 36\%$ vs. $M_{Rep} = 47\%$; $F(1,178) = 13.7, P < 0.001$] compared with Republicans. Conversely, Republicans evaluated a conservative candidate's statements as more true [$M_{Dem} = 44\%$ vs. $M_{Rep} = 54\%$; $F(1,178) = 10.6, P = 0.001$] and less false [$M_{Dem} = 49\%$ vs. $M_{Rep} = 38\%$; $F(1,178) = 12.2, P = 0.001$] compared with Democrats. However, political alignment did not differentially affect flagging rates for either candidate [liberal: $M_{Dem} = 9.8\%$ vs. $M_{Rep} = 10.8\%$; $F(1,178) = 0.29, P = 0.92$; conservative: $M_{Dem} = 8.4\%$ vs. $M_{Rep} = 7.9\%$; $F(1,178) = 0.01, P = 1.00$]. Furthermore, alignment did not interact with social presence to affect "true" responses [liberal: $F(2,178) = 0.40, P = 0.67$; conservative: $F(2,178) = 0.86, P = 0.43$], "false" responses [liberal: $F(2,178) = 0.32, P = 0.73$; conservative: $F(2,178) = 0.99, P = 0.38$], or flagged responses [liberal: $F(2,178) = 0.02, P = 0.98$; conservative: $F(2,178) = 0.004, P = 0.99$]. In sum, although people were more likely to accept a statement as true when its view accorded with their own party affiliation, this alignment neither differentially affected flagging rates nor interacted with condition, distinguishing the effects of social presence from confirmation bias. Perceiving the company of others seemed to influence people's willingness to verify information, not how much they believed it.

Finally, participants rated the candidates as having a high likelihood of lying, regardless of social presence [$M_{Alone} = 5.88, SD = 1.09$; $M_{Group-Small} = 5.55, SD = 1.34$; $M_{Group-Large} = 5.74, SD = 1.37$; $F(2,284) = 1.59, P = 0.21$]. Thus, even in a situation where the

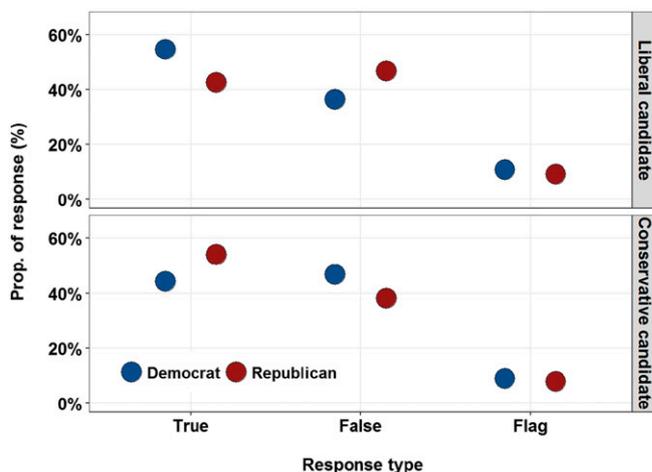


Fig. 2. Proportions of true, false, and flagged (i.e., fact-checked) responses by candidate affiliation and own party identification (experiment 6).

conversational norm of speaking truthfully is violated, people fact-checked less often in group compared with individual settings.

Inducing Vigilance. So far, neither diffusion of responsibility nor conversational norms seem to fully explain why participants fact-checked less often in the company of others. Another possibility proposes that perceiving others' presence might automatically lower people's guards. Some evidence from the regulatory focus literature suggests that individuals with a chronic prevention focus (i.e., those motivated by vigilance) tend to be more accurate with respect to error detection (21). To test the effect of social presence on vigilance, we examined how well people performed on a proofreading task (*SI Appendix, Proofreading*). Participants ($n = 189$; $M_{age} = 36$; 46% female) were asked to identify all errors appearing in a series of passages. For the duration of the task, people saw their own username either displayed on the side of the screen (alone) or alongside 102 others (group).

Those exposed to others found fewer correct errors than those who saw only their names [$M_{Alone} = 10.69, SD = 4.64$; $M_{Group} = 9.77, SD = 5.05$; Wald $\chi^2(1) = 4.16, P = 0.04$], suggesting that social presence may impair vigilance when processing information. To further assess this relationship, the next two experiments tested whether interventions that increase vigilance can promote greater willingness to fact-check ambiguous claims.

Experiment 7.

Method. Previous research has found that others' presence improved performance on a vigilance task when participants believed that these others had access to information about the quality of their performance (22). Experiment 7 investigated the effectiveness of inducing accountability: If people expect their actions to be judged, might they be more receptive to verifying claims? Participants ($n = 330$; $M_{age} = 35$; 42% female) evaluated the same 36 ambiguous headlines from (and under identical incentive specifications as) experiment 1. We assigned respondents to one of three conditions: alone; group (102 others); or group-accountable, where people read that their responses during the task would be revealed to the other 102 participants at the end of the study. Upon evaluating the statements, respondents completed the Regulatory Focus Questionnaire (23) containing two psychometrically distinct subscales on promotion and prevention (1 = never or seldom; 5 = very often). If the failure to fact-check is driven by reduced vigilance in general when processing information, then we would expect individuals who are chronically prevention-focused (a trait associated with being cautious and vigilant) to be less susceptible to social presence effects.

Afterward, participants reported how much they felt accountable for their responses as well as how responsible they felt for

⁵Ninety-four (33%) people identified as Independent, and nine people (3%) identified as other. When all participants are included in the analysis, the effect of social presence on flagging remains, regardless of party identification.

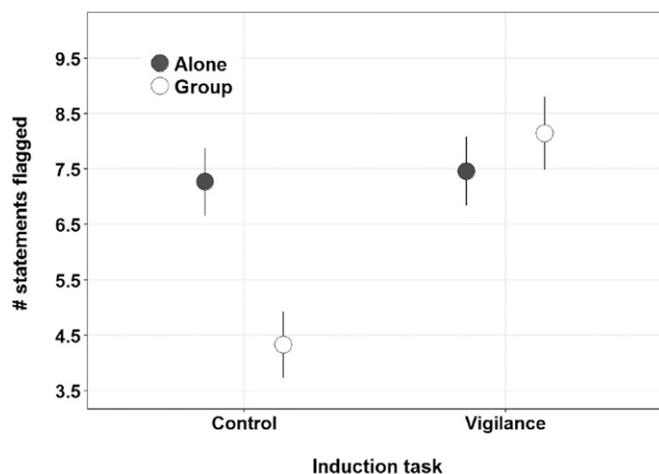


Fig. 3. Number of statements flagged (i.e., amount of fact-checking) by social presence and induction task (experiment 8). Error bars denote SEs (\pm SEM).

discovering the statements' veracity. The participants in the two group conditions further indicated the extent to which they believed that they made their decisions in public (all measures above scaled from 1 = not at all; 7 = very much).

Results. Inducing a sense of public scrutiny indeed made participants in the group-accountable condition feel like they were making their decisions in public more than those in the group condition [$M_{\text{Group}} = 2.10$, $SD = 1.50$; $M_{\text{Group-Accountable}} = 2.90$, $SD = 2.12$; $F(1,327) = 10.2$, $P = 0.001$]. Participants differed along perceived accountability [$F(1,327) = 6.12$, $P = 0.002$]: Those in the group-accountable condition felt equally accountable for their choices as those in the alone condition [$F(1,327) = 0.06$, $P = 0.81$] but more so than those in the group condition [$M_{\text{Alone}} = 4.41$, $SD = 1.75$; $M_{\text{Group}} = 3.58$, $SD = 2.05$; $M_{\text{Group-Accountable}} = 4.34$, $SD = 1.93$; $F(1,327) = 8.52$, $P = 0.004$].

Flagging rates differed overall [$F(2,327) = 3.92$, $P = 0.02$]. Compared with those in the group condition [$M = 4.29$ (12%), $SD = 5.01$], those in the group-accountable condition [$M = 6.75$ (18%), $SD = 7.67$; $F(1,327) = 7.18$, $P = 0.01$] and alone condition [$M = 6.21$ (17%), $SD = 6.86$; $F(1,327) = 4.44$, $P = 0.04$] flagged more statements and did not differ from each other [$F(1,327) = 0.36$, $P = 0.54$]. By generating feelings of accountability through public scrutiny, people may be more inclined to fact-check in the presence of others.

To test the role of individual differences in chronic vigilance on fact-checking behavior, we analyzed flagging rates as a function of the experimental condition (alone, group, and group-accountable) and an indexed measure of prevention scores ($\alpha = 0.86$). We used the following dummy variables: Z_1 given by the contrast comparing the alone condition (-1) with the group condition (+1) and Z_2 given by the contrast comparing the group condition (-1) with the group-accountable condition (+1). An analysis regressing the number of statements fact-checked on Z_1 and Z_2 , prevention scores (z -scored), and their interaction found that chronic prevention focus positively predicted fact-checking [$B = 0.24$, $SE = 0.08$, $t(324) = 3.01$, $P = 0.003$] and did not interact with either Z_1 [$t(324) = -0.19$, $P = 0.85$] or Z_2 [$t(324) = -0.59$, $P = 0.56$]. Individuals who are habitually cautious—being less susceptible to the company of others—tended to fact-check more regardless of social presence.

Experiment 8.

Method. Interpersonal cues, such as accountability to others, aside, can we leverage internal cues to promote fact-checking? Experiment 8 ($n = 385$; $M_{\text{age}} = 34$; 46% female) more directly tested whether enhancing vigilance in the moment can attenuate the resistance toward fact-checking in social settings.

Participants completed three ostensibly unrelated tasks; the first two comprised a vigilance induction shown to heighten prevention focus (24, 25). Specifically, those assigned to the vigilance condition first recalled their past and present duties, obligations, and responsibilities. We then instructed them to avoid missing target words on a subsequent word search puzzle for which they had 2 min to solve. Control participants recalled the layout of their old and current rooms; they were next simply asked to find target words in the same word search. All respondents proceeded to evaluate 36 headlines, per experiment 1, in either an alone or group condition (102 others). **Results.** A 2 (social presence) \times 2 (vigilance) ANOVA on flagged statements found a social presence \times vigilance interaction [$F(1,381) = 8.37$, $P = 0.004$] (Fig. 3). Among those in the control condition, participants again fact-checked fewer statements in the presence of others versus by themselves [$M_{\text{Alone}} = 7.26$ (19%), $SD = 6.62$; $M_{\text{Group}} = 4.33$ (11%), $SD = 4.97$; $F(1,381) = 11.7$, $P = 0.001$]. However, this difference disappeared after exposure to the vigilance induction [$M_{\text{Alone}} = 7.46$ (21%), $SD = 5.69$; $M_{\text{Group}} = 8.14$ (22%), $SD = 7.12$; $F(1,381) = 0.56$, $P = 0.45$]. Although the intervention did not influence people in the alone condition [$F(1,381) = 0.05$, $P = 0.82$], it led those in the group condition to fact-check nearly twice as often [$F(1,381) = 18.3$, $P < 0.001$]. These results further indicate that the asymmetry in fact-checking likely stems from a reduction in vigilance among those in the group condition rather than an increase in vigilance among those in the alone condition.

Summary. Eight experiments furnish convergent evidence that perceiving the presence of others reduces people's willingness to fact-check claims. Using an incentive-compatible real effort paradigm, this pattern held across different payment specifications (lottery-based vs. piece-rate rewards), scoring of flagged responses (gain vs. loss vs. neither), statements (neutral vs. politically charged; less vs. more ambiguous), social contexts (simulated forums with names vs. social media platforms), group sizes (smaller vs. larger), and temporal reference frames (current vs. past others) (*SI Appendix, Table S1*).

We conducted an internal metaanalysis (Fig. 4) that specified a fixed effects model using Hedges' g , a variation of Cohen's d that corrects for biases caused by small sample sizes (26). Limiting our analysis to the differential effect of flagging rates as a function of social presence (i.e., group vs. alone), we observed a medium average estimated effect size of 0.37 ($SE = 0.05$, 95% confidence interval = 0.27, 0.47).

Discussion

Certain features of our paradigm merit discussion. Could participants have simply chosen the flagging option as a proxy for

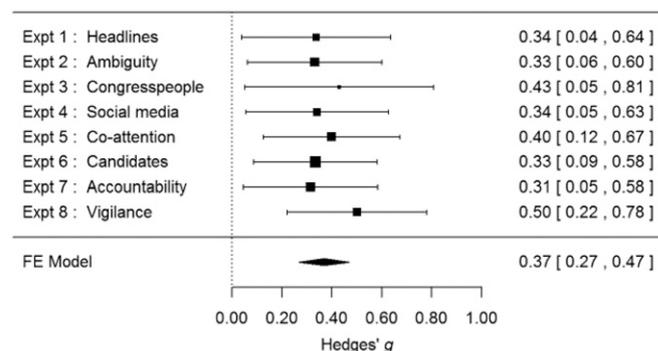


Fig. 4. Internal metaanalysis of social presence using a fixed effects (FE) model (experiments 1–8; $n = 2,238$). Effect sizes (Hedges' g) are derived from comparisons of the number of flagged statements in the alone vs. group conditions. Brackets denote 95% confidence intervals. Sizes of squares are proportional to the precision of the estimate.

uncertainty? People's confidence levels did not vary according to social presence, suggesting that this may not be the case. Another possibility is that evaluating statements alone relative to in the company of others facilitated greater engagement with content. However, those in both the alone and group conditions tended to find the task similarly engaging.

Recall that we raised a few potential reasons why people might be reluctant to fact-check in social settings. The first, diffusion of responsibility, was largely unsubstantiated by the data. In most cases, (i) participants read that their choices were private and confidential (no one reported suspicions to the contrary), and (ii) performance was incentivized on an individual basis. Although a social loafing explanation would predict that self-prominence within a group should eliminate loafing behavior altogether, our participants did not fact-check more when they felt more distinguished and responsible in a group (experiment 3). We observed the same pattern when people evaluated statements presented as Facebook posts without seeing others' names (experiment 4) and when they were exposed to anonymous strangers who had completed the task in the past (experiment 5), further casting doubt on the possibility of free-riding.

Secondly, a conversational norms account would predict that the presence of others heightens the tendency to take information at face value. However, we did not observe a consistent effect of social presence on the number of statements identified as true. Even when people had little reason to find a source particularly credible (experiment 6), they fact-checked less often in a social compared with an individual setting. Furthermore, a separate experiment (*SI Appendix, Following Up*) found that this effect remained when we described fact-checking in more benign terms less likely to violate social conventions—namely, not as flags that appeal to external checks but rather as an option to “follow-up” with the source.

Our data provide some evidence for the third route—reduced vigilance—and suggest that social contexts may impede fact-

checking by, at least in part, lowering people's guards in an almost instinctual fashion. These contexts can take the form of platforms that are inherently social (e.g., Facebook) or can be cued by features of online environments such as “likes” or “shares” that a message receives. These findings, therefore, advance our understanding of how people (mis)interpret information in an increasingly connected world (27).

Taken together, these results lend additional context in the ongoing debate surrounding the spread of false or “fake” news. Note that we have focused on individuals' scrutiny of, rather than belief in, the information that they confront. In reality, a variety of forces likely contribute toward the latter, including confirmation bias, the tendency to persist in one's views despite evidential discrediting, and the allure of sharing controversial or emotionally arousing stimuli (28, 29). Although people may not always be swayed by—or vote according to—the headlines that they see on social media (30), a widespread failure to digest content responsibly nevertheless has the potential to distort both personal beliefs and public opinion. These ideas, after they are entered into the collective consciousness, can prove remarkably “sticky,” even when disproven or invalidated (31).

Spurred by such concerns, recent efforts to promote fact-checking using algorithms and human moderators have found some success in reducing the circulation of unreliable news on social media sites (32). Continuing to devise interventions that encourage greater informational scrutiny poses an important challenge for policymakers who wish to inspire a well-informed populace. Even if we cannot exercise constant vigilance, we would do well to question our wisdom in crowds.

ACKNOWLEDGMENTS. We thank Tory Higgins, Oded Netzer, Jacob Goldenberg, Jaideep Sengupta, Daniel Ames, Don Lehmann, Olivier Toubia, Eva Ascarza, and Ran Kivetz for their feedback. The data in this article are available in the Open Science Framework archives (<https://osf.io/qpfuz/>). This research was supported by the Media and Technology Program at Columbia Business School and the Columbia Business School Research Fund.

- Manjoo F (2016) How the Internet is loosening our grip on the truth. *NY Times*. Available at www.nytimes.com/2016/11/03/technology/how-the-internet-is-loosening-our-grip-on-the-truth.html. Accessed December 30, 2016.
- Pew Research Center (2016) News use across social media platforms 2016. Available at www.journalism.org/2016/05/26/news-use-across-social-media-platforms-2016/. Accessed December 30, 2016.
- Reuters (2016) Digital news report 2016. Available at www.digitalnewsreport.org/. Accessed March 17, 2017.
- Issac M (2016) Facebook mounts effort to limit tide of fake news. *NY Times*. Available at <https://www.nytimes.com/2016/12/15/technology/facebook-fake-news.html>. Accessed March 20, 2017.
- Latané B (1981) The psychology of social impact. *Am Psychol* 36:343–356.
- Argo JJ, Dahl DW, Manchanda RV (2005) The influence of a mere social presence in a retail context. *J Consum Res* 32:207–212.
- Zajonc RB (1965) Social facilitation. *Science* 149:269–274.
- Latané B, Williams K, Harkins S (1979) Many hands make light the work: The causes and consequences of social loafing. *J Pers Soc Psychol* 37:822–832.
- Brickner MA, Harkins SG, Ostrom TM (1986) Effects of personal involvement: Thought-provoking implications for social loafing. *J Pers Soc Psychol* 51:763–769.
- Harkins SG, Jackson JM (1985) The role of evaluation in eliminating social loafing. *Pers Soc Psychol Bull* 11:457–465.
- Garcia SM, Weaver K, Moskowitz GB, Darley JM (2002) Crowded minds: The implicit bystander effect. *J Pers Soc Psychol* 83:843–853.
- Grice HP (1975) Logic and conversation. *Syntax and Semantics, 3: Speech Acts*, eds Cole P, Morgan JL (Academic Press, New York), pp 41–58.
- Clark HH, Shober MF (1992) Asking questions and influencing answers. *Questions About Questions: Inquiries into the Cognitive Bases of Surveys*, ed Tanur JM (Russell Sage Foundation, New York), pp 15–48.
- Schwarz N (2007) Cognitive aspects of survey methodology. *Appl Cogn Psychol* 21:277–287.
- Roberts G (1996) Why individual vigilance declines as group size increases. *Anim Behav* 51:1077–1086.
- Clark RD, III (1974) Risk taking in groups: A social psychological analysis. *J Risk Insur* 41:75–92.
- Petty RE, Cacioppo JT (1986) The elaboration likelihood model of persuasion. *Adv Exp Soc Psychol* 19:123–205.
- Nickerson RS (1998) Confirmation bias: A ubiquitous phenomenon in many guises. *Rev Gen Psychol* 2:175–220.
- van Bommel M, van Prooijen JW, Eijffers H, Van Lange PA (2012) Be aware to care: Public self-awareness leads to a reversal of the bystander effect. *J Exp Soc Psychol* 48:926–930.
- Shteynberg G (2015) Shared Attention. *Perspect Psychol Sci* 10:579–590.
- Förster J, Higgins ET, Bianco AT (2003) Speed/accuracy decisions in task performance: Built-in trade-off or separate strategic concerns? *Organ Behav Hum Decis Process* 90:148–164.
- Klinger E (1969) Feedback effects and social facilitation of vigilance performance: Mere coaction versus potential evaluation. *Psychon Sci* 14:161–162.
- Higgins ET, Shah J, Friedman R (1997) Emotional responses to goal attainment: Strength of regulatory focus as moderator. *J Pers Soc Psychol* 72:515–525.
- Cho CK, Johar GV (2011) Attaining satisfaction. *J Consum Res* 38:622–631.
- Higgins ET, et al. (2001) Achievement orientations from subjective histories of success: Promotion pride versus prevention pride. *Eur J Soc Psychol* 31:3–23.
- Hedges LV, Pigott TD (2001) The power of statistical tests in meta-analysis. *Psychol Methods* 6:203–217.
- Lewandowsky S, Ecker UK, Seifert CM, Schwarz N, Cook J (2012) Misinformation and its correction: Continued influence and successful debiasing. *Psychol Sci Public Interest* 13:106–131.
- Anderson CA, Lepper MR, Ross L (1980) Perseverance of social theories: The role of explanation in the persistence of discredited information. *J Pers Soc Psychol* 39:1037–1049.
- Berger J, Milkman KL (2012) What makes online content viral? *J Mark Res* 49:192–205.
- Allcott H, Gentzkow M (2017) Social media and fake news in the 2016 election. *J Econ Perspect* 31:211–236.
- Gilbert DT, Tafarodi RW, Malone PS (1993) You can't not believe everything you read. *J Pers Soc Psychol* 65:221–233.
- Matias JN (2017) Persuading algorithms with an AI nudge. MIT Media Lab. Available at <https://medium.com/mit-media-lab/persuading-algorithms-with-an-ai-nudge-25c92293df1d>. Accessed March 14, 2017.