Inaccuracy Blindness in Collaboration Persists, even with an Evaluation Prompt

Aimée A. Kane

Palumbo-Donahue School of
Business
Duquesne University
Pittsburgh, USA
kanea@duq.edu

Sara Kiesler

Human Computer Interaction
Institute
Carnegie Mellon University
Pittsburgh, USA
kiesler@cs.cmu.edu

Ruogu Kang

Instagram
Menlo Park, USA
ruogu@instagram.com

ABSTRACT

The tendency to believe and act on others' misinformation is documented in much prior work. This paper focuses on *inaccuracy blindness*, the tendency to take a collaborator's poor information at face value, which reduces problem-solving success. We draw on social psychological research from the 1970s showing that evaluative rating scales can prompt a change in perspective. In a series of studies, we prototyped and tested an evaluation prompt meant to encourage skepticism in participant detectives trying to identify a serial killer. In tests of the prototype, the prompt was partially successful in inducing skepticism (Exp. 1), but a larger study (Exp. 2) showed that, despite the evaluation prompt, participants' inaccuracy blindness persisted. This work, and the literature more generally, shows that the tendency to be misled by collaborators' inaccurate information remains a strong phenomenon that is hard to counteract and remains a significant challenge for the CHI community.

Author Keywords

CSCW; distributed work; empirical studies; collaboration; communication; social influence; misinformation; groups; problem solving; information sharing; sensemaking

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Our society is experiencing a serious increase in misinformation that is spread online from person to person. Take a message about aspirins and heart attacks falsely attributed to a Mayo Clinic cardiologist that began circulating among acquaintances in 2009. Despite being

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CHI 2018, April 21–26, 2018, Montreal, QC, Canada
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-5620-6/18/04...\$15.00
<https://doi.org/10.1145/3173574.3174068>

Victim Info Zone 8 ☐ Current ☒ Cold

Incident type
☐ Accident (self-inflicted)
☐ Accident (inflicted by other)
☐ Suicide or intentional self-harm
☒ Interpersonal violence-related (incl homicide)
☐ Legal intervention (injured by Law enforcement)

Incident Location Type 1
 1-House/apartment
 2-Street/road, sidewalk, alley
 3-Highway/freeway
 4-Bar, nightclub
 5-Bank, ATM
 6-Liquor store
 7-Other commercial estab.
 8-Office building
 9-Parking lot/public garage
 10-Sports/athletic arena
 11-School
 12-College/University
 13-Public transportation
 14-Hospital/medical facility
 15-Farm
 16-Jail, prison, detention center
 17-Park, playground, public use ar
 18-Natural area (field, woods)
 19-Other

Manner of Death (from ME/C report)
☐ natural ☒ homicide
☐ accident ☐ pending
☐ suicide ☐ undetermined

Death Date: 5/23/10
 Death Time: 4:00 – 5:00 PM

Homeless ☒ No ☐ Yes ☐ Unknown
 Marital Status ☐ Married ☐ Divorced ☐ Never Married ☐ Unknown ☐ Widowed

Person Info				
Person Name	Person Type	Age	Sex	Race
1. Vickie Fisk	Victim	19	F	cauc
2. Julia Tutt	Roommate	20	F	cauc
3. Dennis Rathbone	Boyfriend	21	M	cauc
4. Roger McClure	Landlord	38	M	cauc
5. Sean Ludwig	Neighbor	34	M	cauc

Figure 1. Excerpt of a case file examined by participants, acting as detectives.

debunked by the Mayo Clinic, three years later it was still being passed on by ordinary people to friends with cardiac ailments, garnering more than 33,000 shares on Facebook [https://www.huffingtonpost.com/2013/09/24/online-myths_n_3954799.html].

Many researchers, including those within the CHI community, are playing a helpful role in understanding the mechanisms that explain the spread of misinformation and in designing systems that might identify, counteract, and mitigate the phenomenon [5], and more generally to explore, test, and report improvements to design, policy, or other remedies to support decision making, e.g., [23].

This paper focuses on one aspect of the larger misinformation phenomenon, the tendency for collaborators to take at face value misinformation from a partner, called *inaccuracy blindness*. We report on prototyping of a design feature meant to prompt skepticism of a partner's information and two experiments to test its effectiveness. We used a collaborative analysis task in which participants, playing the role of a police detective, needed to sort through crime case files (see Figure 1) to identify a serial killer.

Each participant received a report purportedly written by a detective partner who had been analyzing the files, and we experimentally varied the accuracy of the information in the partner's progress report. In the prototyping of the evaluative rating prompt, we found some evidence of the desired change in perspective. In Experiment 1, the evaluation prompt was partially successful in inducing skepticism of partner's inaccurate information. However, in Experiment 2, a larger study that compared accurate with inaccurate partner information, inaccuracy blindness persisted, despite the evaluation prompt. These studies, and the literature more broadly, suggest that the tendency to be misled by collaborators' misinformation is arduous to counteract and a significant challenge for the CHI community.

Scope of the problem

Misinformation transmitted from person to person no doubt has always plagued humans. Research on the adoption of false information tells us that people in groups, including collaborations, are especially likely to be misled by others. This work goes back to the famous Asch studies [1]. Studies of "groupthink" pointed to the tendency of group members to rely too much on the opinions of other group members [10]. With advances in communications and information technology, and the advent of distributed work, the misinformation problem has seemingly become much larger. Real-world examples include people's adoption of bad online advice on financial investments or health [4, 21]. Recent experiments show that the presence of purported co-workers reduces fact checking [11], and a series of studies confirms that misinformation from a collaborator or advisor predicts poor problem solving [13] or decision making [15]. This work suggests that the power of misinformation is made all the greater if its source is social—other persons.

Theoretical background

Researchers have studied why misinformation from people has a powerful influence, despite possible or even probable negative consequences. The causes include lack of expertise, time, or distraction, e.g., [22], lack of attention or vigilance [11], a tendency to prefer causal reasoning over denials [3], and a reliance on heuristics, including social cues, e.g., [14]. People even believe misinformation in the face of factual corrections [16], and they will say misinformation is "valuable" when it has clearly harmed their decisions [13]. In short, a multiplicity of cognitive and social phenomena lead people to take what others say at face value.

Researchers have explored possible interface design features that help prevent, or at least, mitigate inaccuracy blindness [2]. Some studies [7, 13] build on the concept of "sensemaking" to investigate if various sensemaking tools, such as annotators and progress reports, aid problem solving. These studies show some promise, but tools have limited value if people overly-rely on poor information from others. What might interrupt people's almost

automatic reliance on the opinions of others? There is some research in the CSCW literature suggesting that displaying controversy or inconsistency in others might interrupt inaccuracy blindness [19], Exp. 2 in [13]. However, for controversy or inconsistency in a group to be useful, it would need to be noticed. Prior research suggests that in situations where people are called upon to evaluate another's information before they can return to their task, they will more thoroughly consider the good or poor quality of the other's ideas [12]. The next step would be to induce this skepticism, but without also making people doubt the value of the collaboration itself.

We drew on early research by Thomas Ostrom and colleagues suggesting that people can be prompted or "primed" [6, 17] to consider the possibility that information from a collaborator might be misleading by the use of a simple evaluation rating form that ranges from bad to good. The theoretical argument is that just seeing a scale that ranges from bad to good will lead people to consider that the information they receive could be of bad or good quality.

Prototyping

Building on the prior work cited above, we decided to incorporate an evaluation rating scale that would be seen by participants when they saw a partner's progress report and prior to their own analysis. We started with a scale that read as follows:

Please read the partner's report and rate the extent to which you think the report could be helpful in solving the homicide cases. The rating scale goes from 5 (Extremely Helpful) to 1 (Extremely Misleading).

Extremely Helpful	Somewhat Helpful	Neither Helpful Nor Misleading	Somewhat Misleading	Extremely Misleading
5	4	3	2	1

We conducted think aloud interviews with 5 participants who were asked to read an inaccurate partner report and then evaluate it using this scale. Most indicated confusion, saying that it was challenging to rate the report using this scale because they did not have access to the case files and therefore did not know whether the report would be helpful or misleading. Some participants suggested, however, that they could rate the logic of the report.

We therefore created the following evaluative scale, which, like the previous one, displays a range of options suggesting another's information may range from good to poor quality.

Please read the partner's report and rate how logical the report seems to you. The rating scale goes from 5 (Extremely Logical) to 1 (Not Logical).

Entirely Logical	Mostly Logical	Somewhat Logical	A little Logical	Not Logical
5	4	3	2	1

You (Pat) are a [redacted] University student working as a detective intern with the [redacted] University police department. Your partner, Alex, is a [redacted] student working as a detective intern with the [redacted] University police department. Your partnership has been assigned 6 unsolved homicides in the [redacted] area. Your task is to solve as many of these cases as possible, which involves identifying suspects and the evidence to support conviction (e.g., clues).

During Phase 1 of the experiment, your detective partner, Alex, worked on analyzing these homicide cases and then wrote a progress report for you.

Now, during Phase 2 of the experiment, you are taking over the task, and continuing the analysis that was begun by your partner and is summarized in the progress report shown below. Please read your partner's report now and rate how logical the report seems to you.

Show report :: Show cases

Alex's Progress Report
 [redacted] student detective

Darleen Raffield - While husband had access to poison, he had no motive. Most likely suspects are Arturo Talamo and her boss Wade McMonagle. Boss and Arturo have both previously threatened the victim. Arturo managed financials while Boss really didn't like victim's new job. Boss had access by buying the victim a drink, who left shortly after, allowing enough time for the poison to circulate and kill her by the time the body was found.

Madeleine Moeller - Sylvia Goad saw a man carrying a small toolbox follow her off the bus. Her husband Frank Moeller may have done it, because he is a landscaper and has access to tools. He could have taken off work or left work early, and he also had access to their home.

Vickie Fisk - The prime suspect is Roger McClure, the landlord. He had prior record. The victim had complained to the landlord about front door trouble before. Lock was already broken. Landlord knew about it so has opportunity.

Ruby Forman - Victim died at her own home. There wasn't any signs of break-ins. Could be done by someone that the victim knows. Most likely, it is her new boyfriend, George Dugger. As the victim's ex-husband was out of town, he was ruled out to be a suspect in this case.

Click here to write the final report:

Write your own report

Please rate how logical Alex's report is:

☐ 1. Not logical at all
 ☐ 2. Not logical
 ☐ 3. Neither logical nor illogical
 ☐ 4. Logical
 ☐ 5. Very logical

submit

Figure 2. Experiment 1. Screen shot showing the presentation of inaccurate information in partner's ("Alex") progress report and the evaluation prompt to rate the logic of the progress report. (We have redacted organizational information.)

Eight participants reported no confusion or inability to use the scale to rate another's progress report. Examples: "[The] rating scale helped know what to pay attention to," and "I looked at the rating scale and focused on how logical the report was."

With this feedback, we decided to test the effectiveness of the evaluation prompt experimentally. We hypothesized:

Prompting participants to rate the logic of a partner's progress report will increase participants' skepticism of inaccurate information from the partner, and will increase their problem solving success.

EXPERIMENT 1

As noted above, participants played the role of a police detective who had to examine bus routes, crime statistics, police reports, and witness accounts to identify a serial killer. We adapted the scenario, problem task, and procedure from prior work by Balakrishnan et al. [2] and Kang et al. [13].

Method

Participants and design

We recruited 79 undergraduate students who received extra credit in a business course for participating. We applied two related *a priori* inclusion criteria to ensure data quality and

consistency with prior work. Because the case records were written in English and contained Western police jargon as well as idiomatic language from witness reports, we restricted participants to native English speakers. We also did not use participants incapable of solving the two practice cases. We thus excluded 7 participants for a sample of 72.

Procedure

The participants were told they should try to solve a serial killer case by correctly identifying the serial killer from an array of possible suspects contained in case materials from previous murders. Participants reviewed digital case files and other data such as bus route information and maps showing the locations of various murders. To solve the task, participants needed to piece together clues that identified the serial killer responsible for four of the six homicides described in the case data.

Participants worked alone and they were first trained in two somewhat easier but analogous detective problems. Then, just prior to working on the serial killer task, they were given a webpage on which they could read a progress report from a purported detective partner who had been working on the same serial killer case (Figure 2). (Misleading information was amalgamated from actual participants'

Partner is from:	Intervention Designed to Encourage Skepticism	
	Evaluation Prompt	No Evaluation Prompt
Same University	18	17
Different University	19	18

Table 1. Design of Experiment 1.

inaccurate reports obtained in previous studies using the serial killer paradigm, and by permission of the authors.) The inaccuracies consisted of misleading clues, such as witness reports pointing to incorrect suspects. We examined whether participants were able to solve the serial killer case, despite all receiving the same misleading information provided by the partner.

As displayed in Table 1, we varied whether or not participants were given an evaluation prompt to rate the partner's progress report, and half the participants were told the partner was from the same university and half were told the partner was from a different university. Because partner organization does not change the main findings, we reserve discussion of this variable to later in the paper.

Results of Experiment 1

Participants spent significantly more time (measured in seconds) reading their partner's report when prompted to evaluate its logic, $M = 224.9$, $SD = 44.9$, than did their counterparts who received no evaluation prompt, $M = 180.2$, $SD = 44.9$, $t(69) = 4.82$, $p < .001$. As shown in Figure 3, the evaluation prompt had a modest effect on participants' ability in their final report to identify the true serial killer rather than the wrong suspect or no suspect at

Figure 3. Experiment 1 solve rate for participants prompted (or not) to evaluate a report containing inaccurate information. Error bars represent the standard errors of the means.

all. More participants solved the case when they had been prompted to evaluate the logic of a misinformed partner report, 43% ($SE = .08$), than when they had received no such evaluation prompt, 31% ($SE = .08$), but this was not a statistically significant difference, $\chi^2(1) = 1.1$, $p = .30$.

To delve into how the prompt worked, or failed to work, see Figures 4 and 5, which tell a story of how the prompt changed reactions to the partner's misinformation. What is particularly noticeable is that only 28% of participants, 95% confidence interval (CI) [.17, .44], rated the partner's misinformed report at the middle or bottom of the scale, and 72% of the participants, 95% CI [.56, .84], in spite of seeing the misinformed report, rated it as "logical" or "very logical" (Figure 4). As shown in Figure 5, an encouraging finding was that higher solve rates are associated with

Figure 4. How logical or illogical the partner's report was rated when it was actually inaccurate and participants were prompted to evaluate it.

Figure 5. Experiment 1 solve rate as a function of participants' evaluative ratings of inaccurate information from a partner. Error bars represent the standard errors of the means.

rating the report as “not logical at all” or “not logical” (i.e., 100% and 50%, respectively), than with rating the report as neutral, logical, or very logical (i.e., 33%, 43%, and 33%, respectively). (To examine this finding statistically, more participants would be needed in the less logical categories.)

Discussion of Experiment 1.

The evaluation prompt increased the time participants spent reviewing a partner’s report and had some modest success in raising the solve rates. Nonetheless, there was a tendency for people to take inaccurate information at face value, as shown with the evaluative ratings that cluster at the logical end of the scale in Figure 4. These positive evaluations of misinformation can harm problem-solving judgments as shown in the right side of Figure 5. A non-negligible set of people do not take the inaccurate information at face value, instead rating it at the less logical end of the scale (see Figure 4), and these individuals are relatively likely to solve the case (see Figure 5). Said another way, the evaluation prompt appears to help remediate misinformation, but only with people who rate the report as less logical. These findings imply that we should design to prompt attention to a collaborator’s information quality — skepticism of inaccurate information but also recognition of accurate information. Experiment 2 follows this argument, examining whether an evaluation prompt helps collaborators become more skeptical, without undermining their ability to learn from partners.

EXPERIMENT 2

The method and procedure of this study was similar to that of Experiment 1. However, participants received either the inaccurate report used in Experiment 1, or a report of equal length containing accurate information. (The accurate information was amalgamated from participants’ progress

reports obtained in previous studies, and by permission of the authors.) All participants received the evaluation prompt. The goal was to understand how well participants receiving inaccurate information would do, versus how well they would do with an accurate report from the partner, and if they would be able to overcome misleading information from the partner to solve the serial killer case. Applying the same inclusion criteria as in Experiment 1, the total N was 138, randomly assigned to each condition, 67 in the accurate information condition and 71 in the inaccurate information condition.

Results of Experiment 2.

Participants spent significantly more time (measured in seconds) reading their partner’s report when it contained inaccurate information, $M = 222.0$, $SD = 53.5$, than when it contained accurate information, $M = 202.8$, $SD = 44.1$, $t(136) = 2.30$, $p = .01$.

Overall, accurate information from the partner was helpful, raising the solve rate to 73% of participants, ($SE = .05$), which is somewhat better than the approximately 50% solve rates in prior studies of pairs of actual collaborating participants in real time and is about the same as when collaborators have additional interface tools [2, 7]. Solve rates for participants who received an inaccurate report dipped to 39% ($SE = .06$), (see Figure 6), which represents significantly poorer performance, $\chi^2(1) = 15.2$, $p < .001$. This finding suggests that the evaluation prompt meant to initiate skepticism did not overcome inaccuracy blindness.

We then delved into the relationship of the actual evaluation prompt ratings and success rates. Figure 7 shows

Figure 6. Experiment 2 solve rate for participants prompted to evaluate the logic of a partner’s report containing inaccurate vs. accurate information. Error bars represent the standard errors of the means.

Figure 7. Experiment 2 how logical or illogical the partner’s report was rated when it was actually either inaccurate vs. accurate, and participants were prompted to evaluate it.

Figure 8. Experiment 2 solve rate as a function of participants' evaluative ratings when receiving either inaccurate or accurate information from a partner. Error bars represent the standard errors of the means.

that participants rated a partner's reports as more logical when it contained accurate information, $M = 4.06$, $SD = .69$, than when it contained inaccurate information, $M = 3.63$, $SD = .76$, $t(136) = 3.43$, $p < .001$. Although more participants used the middle or the bottom of the logic scale when exposed to inaccurate information, 30% with 95% CI [.20, .41] than when exposed to accurate information, 12% with 95% CI [.06, .22], few participants rated information using these less logical categories.

The findings shown in Figure 8, suggest that solve rates depended on the accuracy of the partner's information and whether participants were actually prompted to be skeptical (i.e., rated the partner's information as illogical). The single participant who rated the partner's report "1. Not logical at all," solved the case, despite receiving misinformation. Among participants who rated the report "2. Not logical," 33% solved the case, regardless of whether the partner's information was accurate or inaccurate. This finding suggests a potentially harmful consequence of skepticism. Among participants who rated the report more highly (i.e., neutral, logical, or very logical), those receiving accurate information were much more likely to solve the case (60%, 73%, and 87%, respectively) than those receiving inaccurate information (43%, 38%, and 33%, respectively). These last results show that skepticism did not overcome misinformation; neither did it harm collaboration when information was deemed good.

Role of partner's organization

When we began this work, we had speculated that, if the partner was from the same organization (i.e., the same university) as the participant, that he or she might react differently to the partner's report than if the partner were an outsider, that is, from a different organization. A relevant theory, social identity theory, pertains to the identity of the

source of the misinformation; people sometimes display adherence to erroneous ideas from insiders, and concomitantly skepticism of erroneous ideas presented by outsiders [20]. To the extent this work might generalize to private problem-solving settings, it suggests that people may think twice before adopting misinformation from someone who belongs to a different organization, and that skepticism could reduce their adoption of the other's misleading information. We tested this idea experimentally by varying the home university of the participants' partner. In Experiment 1, this was an asymmetric manipulation, but in Experiment 2, participants were actually from two universities, and we manipulated the partner's home university.

The findings did not replicate previous work, possibly because group identity was not emphasized and was not a strong contextual factor in these studies. In Experiment 1, participants tended to be (but not statistically significantly) more skeptical of partners from their own university, and in Experiment 2, the partner's organization did not matter and did not change the main findings.

Role of participant's ability

The studies provide some evidence that an evaluation prompt might raise participants' skepticism of misinformation (observable in low evaluative ratings), which could improve problem solving (see Figures 5 and 8). A reviewer suggested there may be a third factor: participants with stronger ability could evaluate inaccurate reports as less logical and also perform well. To look into this possibility, we checked whether participants' initial ability (i.e., approximated as performance on the two training cases given before they ever saw their partner's report) predicted ratings of partners' inaccurate reports. We found that practice case performance was not correlated with evaluative ratings in Experiment 1, $r = -.04$, $p = .84$ or in Experiment 2, $r = .06$, $p = .61$. These correlations are not proof, but they tend to show that initial ability does not explain the evaluative ratings.

GENERAL DISCUSSION

To review this work:

First, we designed and prototyped a prompt whose intent was to induce skepticism of an inaccurate report by a collaborative partner. The design of the prompt built on early research, e.g., [6], suggesting that skepticism of another's report might increase if people had to rate the report on a scale with options ranging from illogical to logical.

Second, we tested the prototype, comparing how the presence of the prompt (vs. no prompt) changed perceptions and solve rates (Exp. 1). The prompt had some influence, but only with people who actually rated the report as less logical.

Third, we conducted a replication in which we compared participants' perceptions and solve rates when they used the

prompt to evaluate accurate vs. inaccurate reports by the partner (Exp. 2). Again, solve rates increased but only with the few who rated the inaccurate report as more illogical. Overall, solve rates were low when the partner had an inaccurate report.

The conclusion we reach from these findings is our title: inaccuracy blindness in collaboration persists, even with an evaluation prompt. This finding is not unique in the CHI literature. Zhang, Bellamy, and Kellogg [23], similarly, found that two “debiasing” remedies they explored to reduce heuristic processing and encourage better decisions were only partly effective.

Are there remedies for inaccuracy blindness?

Inaccuracy blindness is a specific instance of a general problem in our society of the transmission of misinformation. Two general strategies characterize attempted remedies so far. One strategy consists of computationally identifying misinformation so people can avoid it. Some promising approaches combine behavioral data, crowdsourcing, comparisons of news sites, linguistic cue learning, and network analysis [5]. This strategy, even though promising, has major limitations when applied to inaccurate information within collaborations, or even within organizations. The data that could be used for computational analysis is likely to be comparatively small and somewhat unique (e.g., about a specific crime with limited case information and only a few collaborators). Automated analyses based on limited unique problems risk error and false positives. More importantly, even if the software correctly tagged information as potentially false, such as, “This report may be inaccurate,” our current study suggests collaborators might not attend to such a message.

Another strategy has been to develop applications and tools that help collaborators do a better job of analysis. For instance, tools, such as visualizations of the connections among suspects in a crime, can be powerful aids in problem solving and decision-making. Balakrishnan et al., using the same general experimental paradigm as ours with synchronous real collaborators, raised solution rates to 60% [2] and Goyal and Fussell raised rates to approximately 75% [7]. Neither study, however, was designed to identify or test the effect of misinformation. Unfortunately, the impact of collaborators’ or co-workers false information can undermine even strong aids to correct analysis because human nature is so strongly attuned to social cues, the presence of others, and the connecting tissue created by collaboration, e.g., [11]. Thus, we may need a combination of tools, including more powerful notifications that interrupt the powerful social heuristic — taking what others say at face value.

The goal of a warning about misinformation in collaborations must have two key attributes: not to undermine collaborators from adopting accurate information from partners (“do no harm”), and at the same

time to cause collaborators to exert discriminatory thinking and analysis that could help them reject inaccurate information. These requirements will make design difficult, as they did in our case. The balance we struck in the present study, of having collaborators rate another’s work on an illogical to logical scale, was apparently too weak. Prior work by Starbird et al. [18] suggests a stronger tactic. These authors found that overt expressions of uncertainty tend to lead up to denials of the truth of rumors. This finding suggests we might try inducing uncertainty with prompts, such as “Please reflect on whether the report could be wrong,” or asking for active reflection, “Why might this report be inaccurate?” This tactic might also reduce the heuristic thinking that is involved with social cues and social presence of others, but it does risk reducing trust.

A more social remedy to people’s tendency to accept others’ misinformation might be to encourage them to check their first impressions or their own analyses with other members of their collaboration. Prior work, such as Greis et al. [8] suggests that people actually prefer having multiple opinions and that multiple opinions give them confidence. Multiple opinions could improve problem solving and decision making if they made good information apparent or encouraged deeper processing. For instance, Kang et al. [13] found that mixing one inaccurate with one accurate report improved problem solving. The difficulty here would be to insure (a) people seek out other opinions, and (b) that accurate information was included in what they actually received. It would be important, as well, not to interfere with trust and efficiency in the collaboration.

Limitations and Future Work

The limitations of this work should be noted. First, although our intention was to replicate the circumstances in which inaccuracy blindness can occur, our participants were not trained, experienced detectives who might have picked up on their partner’s misleading information. Another limitation is that the case materials were a combination of text materials and maps and lacked tools that might have improved success rates [7]. We also focused here on small rather than large collaborations. Results might differ in open settings where total strangers provide misinformation as such settings may tend to induce skepticism. There is much room for studying experts, misinformation with stronger tools, and larger, open collaborations.

CONCLUSION

Online collaborations and workgroups that include strangers, revolving membership, partners or advisors with social credentials but perhaps unknown expertise, and work done over distributed time and space, are more prevalent than ever. These interactions in today’s world are unavoidable and their value is undeniable, but the risks of having collaborators who provide sloppy, false, misleading, or inaccurate information have increased as well.

Although there are documented real world cases of the impact of poor collaboration (including collaboration failures that did not prevent the 911 attacks), and prior work has called out the problem of misinformation from others [9, 19] and identified the existence of inaccuracy blindness in small collaborations [13], the current work advances our understanding of inaccuracy blindness and contributes to the CHI community a new idea for remediation and controlled experimentation on the remedy. This work demonstrates inaccuracy blindness persisted despite a prompt that went well beyond a simple warning—it required the participants to attend to and evaluate their partner's information. Thus we document the strength of inaccuracy blindness in collaborations and challenge our community to address this problem. Doing so at the level required might risk undermining collaboration ties, so it is a difficult challenge but would be one important step in solving the larger misinformation problem faced by our society.

ACKNOWLEDGMENTS

We especially thank Elise Carter, Sarah Deluiis, and Lisa Kim, the undergraduate assistants who helped with the experiments, and gratefully acknowledge grants from the National Science Foundation (IIS-1026234, IIS-1129642, and OCI-1025656). We also thank our reviewers for their detailed and incisive suggestions for further analyses and discussion of the results.

REFERENCES

1. Solomon E. Asch. 1956. Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psych. Monographs* 70, 9: Whole No. 416.
2. Aruna D. Balakrishnan, Susan R. Fussell, Sara Kiesler. 2008. Do visualizations improve synchronous remote collaboration?. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '08), 1227-1236. <https://doi.org/10.1145/1357054.1357246>
3. Adam Berensky. 2015. Rumors and health care reform: experiments in political misinformation. *British J. of Pol. Sci.* 47: 241-262.
4. Ann Brenoff. 2013. What bad advice did this heart attack survivor get? *The Huffington Post*. Retrieved September 7, 2017 from http://www.huffingtonpost.com/2013/09/24/online-myths_n_3954799.html
5. Niall J. Controy, Victoria L. Rubin, Yimin Chen. 2015. Automatic deception detection: methods for finding fake news. In *Proceedings of the Association for Information Science and Technology*, 52, 1-4. <https://doi.org/10.1002/pra2.2015.145052010082>
6. Katherine M. Gannon and Thomas M. Ostrom. 1996. How meaning is given to rating scales: the effects of response language on category activation. *J. Exp. Soc. Psych.* 32, 4: 337-360.
7. Nitesh Goyal and Susan R. Fussell. 2016. Effects of sensemaking translucence on distributed collaborative analysis. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing* (CSCW '16), 288-302. <https://doi.org/10.1145/2818048.2820071>
8. Miriam Greis, Emre Avci, Albrecht Schmidt, Tonja Machulla. 2017. Increasing users' confidence in uncertain data by aggregating data from multiple sources. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '17), 828-840. <https://doi.org/10.1145/3025453.3025998>
9. Derek L. Hansen, Patrick J. Schone, Douglas Corey, Matthew Reid, Jake Gehring. 2013. Quality control mechanisms for crowdsourcing: peer review, arbitration, & expertise at familysearch indexing. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work* (CSCW '13), 649-660. <https://doi.org/10.1145/2441776.2441848>
10. Irving I. Janis. 1972. *Victims of groupthink: A psychological study of foreign-policy decisions and fiascoes*. Houghton Mifflin.
11. Youjung Jun, Rachel Meng, Gita Venkataramani Johar. 2017. Perceived social presence reduces fact-checking. *PNAS* 114, 23: 5976-5981.
12. Aimee A Kane. 2010. Unlocking knowledge transfer potential: knowledge demonstrability and superordinate social identity. *Organization Sci.* 21, 3: 643-660.
13. Ruogu Kang, Aimee A. Kane, Sara Kiesler. 2014. Teammate inaccuracy blindness: when information sharing tools hinder collaborative analysis. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing* (CSCW '14), 797-806. <http://dx.doi.org/10.1145/2531602.2531681>
14. Durairaj Maheswaran and Shelly Chaiken. 1991. Promoting systematic processing in low-motivation settings: effect of incongruent information on processing and judgment. *J. Pers. Soc. Psych.* 61,1: 13-25.
15. Duyen T. Nguyen, Laura A. Dabbish, Sara Kiesler. 2015. The perverse effects of social transparency on online advice taking. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work & Social Computing* (CSCW '15), 207-217. <http://dx.doi.org/10.1145/2675133.2675253>
16. Brendan Nyhan and Jason Reifler. 2015. Displacing misinformation about events: an experimental test of causal corrections. *Exp. Pol. Sci.* 2: 81-93.
17. Thomas M. Ostrom. 1970. Perspective as a determinant of attitude change. *J. Exp. Soc. Psych.* 6: 3: 280-292.

18. Kate Starbird, Emma Spiro, Isabelle Edwards, Kaitlyn Zhou, Jim Maddock, Sindu Narasimhan. 2016. Could this be true? I think so! Expressed uncertainty in online rumoring. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '08), 1227-1236.
<http://dx.doi.org/10.1145/2858036.2858551>
19. W. Ben Towne, Aniket Kittur, Peter Kinnaird, James Herbsleb. 2013. Your process is showing: controversy management and perceived quality in wikipedia. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work* (CSCW '13), 1059-1068.
<http://dx.doi.org/10.1145/2441776.2441896>
20. Henri Tajfel and John C. Turner. 1979. An integrative theory of group conflict. In *The Social Psychology of Intergroup Relations*, William G. Austin and Stephen Worchel (eds.). Brooks/Cole, Monterey, CA, 33-47.
21. James Watkins. 2014. Online investment advice and robo-advisers. Retrieved August 6, 2017 from https://www.huffingtonpost.com/paladinregistrycom/online-investment-advice-and-robo-advisers_b_5606918.html
22. Wendy Wood, Nancy Rhodes, Michael Biek. 1995. Working knowledge and attitude strength: an information-processing analysis. In *Attitude Strength: Antecedents and Consequences*, Richard Petty & Jon Krosnik (eds.). Erlbaum, Mahwah, NJ, 283-313.
23. Yunfeng Zhang, Rachel K.E. Bellamy, Wendy A. Kellogg. 2015. Designing information for remediating cognitive biases in decision-making. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '15), 2211-2220.
<http://dx.doi.org/10.1145/2702123.2702239>