

Associative Inference Can Increase People’s Susceptibility to Misinformation

Sian Lee, Haeseung Seo, Dongwon Lee, Aiping Xiong

The Pennsylvania State University, USA
{szl43, hxs378, dongwon, axx29}@psu.edu

Abstract

Associative inference is an adaptive, constructive process of memory that allows people to link related information to make novel connections. We conducted three online human-subjects experiments investigating participants’ susceptibility to associatively inferred misinformation and its interaction with their cognitive ability and how news articles were presented. In each experiment, participants completed recognition and perceived accuracy rating tasks for the snippets of news articles in a tweet format across two phases. At Phase 1, participants viewed real news only. At Phase 2, participants viewed both real and fake news. Critically, we varied whether the fake news at Phase 2 was inferred from (i.e., associative inference), associated with (i.e., association only), or irrelevant to (i.e., control) the corresponding real news pairs at Phase 1. Both recognition and perceived accuracy results showed that participants in the associative inference condition were more susceptible to fake news than those in the other conditions. Furthermore, hashtags embedded within the tweets made the obtained effects evident only for the participants of higher cognitive ability. Our findings reveal that associative inference can be a basis for individuals’ susceptibility to misinformation, especially for those of higher cognitive ability. We conclude by discussing the implications of our results for understanding and mitigating misinformation on social media platforms.

Introduction

The ubiquitousness of social media platforms and people’s extended use of them for news consumption have partially helped bring a proliferation of misinformation (Allcott and Gentzkow 2017; Lazer et al. 2018). The extensive spread of fake news on social media platforms (Vosoughi, Roy, and Aral 2018) can have serious negative impacts on individuals and society, including election manipulation (Rosenberg 2016), spreading false treatments of COVID-19 (Perlow 2020), as well as reducing people’s COVID-19 vaccination intention (Loomba et al. 2021).

Fake news especially becomes “successful” when it could cast doubts on truth or deceive news consumers. Detecting such fake news and preventing its harm, thus, have deep intellectual values as well as broad societal impacts. To help mitigate the negative effects of fake news on social media

platforms, we seek to understand **cognitive processes** that increase people’s susceptibility to *remember* and *believe in* misinformation.

Prior studies have examined various cognitive factors impacting people’s *false memory* of fake news (Loftus 2005). Research has shown that repeated exposure can increase people’s familiarity with misinformation and thus enhance their recollection afterward (Foster et al. 2012). Compared to individuals with higher cognitive ability, those with lower cognitive ability were more likely to remember fake news in agreement with their political stance (Murphy et al. 2019).

The impacts of repetition (Begg, Anas, and Farinacci 1992) and cognitive ability (Bago, Rand, and Pennycook 2020) have also been evident in people’s *belief* in fake news. People increased their misbelief even if the repeated fake headlines were disputed (Pennycook, Cannon, and Rand 2018; Seo, Xiong, and Lee 2019). Individuals’ cognitive ability levels were negatively correlated with the accuracy ratings of fake news but positively correlated with the accuracy ratings of real news (Pennycook and Rand 2019).

We examine the impact of **associative inference**, an adaptive, constructive memory process that allows people to link together extracted knowledge (Zeithamova and Preston 2010) or acquired experience (Carpenter and Schacter 2017) that shares a common feature to make novel connections. When one person views *false inferences* based on associated real news articles (see Figure 1), feelings of familiarity comparable to or stronger than repetition could be created, increasing her or his susceptibility. Different from a piece of fake news that is completely fabricated, associatively inferred misinformation represents a dedicated attempt for deception by leveraging people’s prior knowledge of existing real news.

Few studies have been conducted on people’s susceptibility to associatively inferred fake news (Lee et al. 2020; Xiong et al. 2023). While Xiong et al. found that associative inference could increase participants’ susceptibility to misinformation in the familiarity judgment (i.e., recognition), the existing studies do not provide conclusive results on participants’ misbelief evaluation. We fill the gap in this work.

In addition, using hashtags, prefixed by the # symbol with a keyword or keywords, has become a common tagging method to help social media users associate tweet messages and share certain events or contexts (Bonilla and Rosa

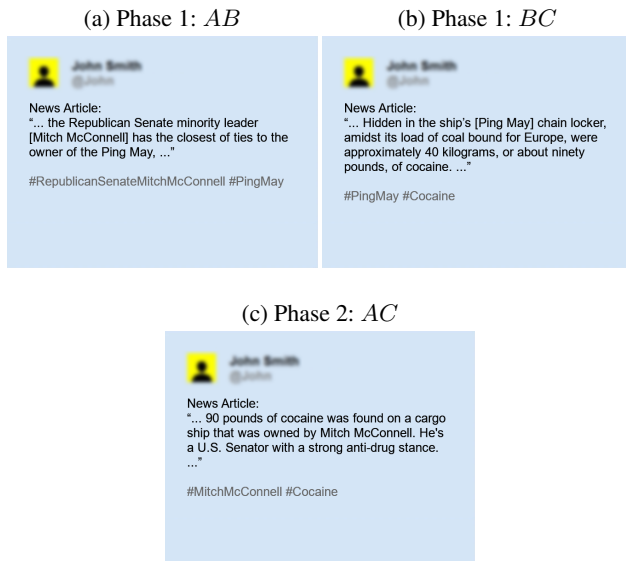


Figure 1: An associatively-inferred triad ($AB&BC \rightarrow AC$) of Experiments 1. The top row shows an overlapping, real-news pair ($AB&BC$) of Phase 1, each of which is associated with a piece of fake news (AC) of Phase 2 (the bottom row) through a keyword listed as a hashtag.

2015). Potentially, usage of two or more hashtags visually enhances the parts following # (see Figure 1), which may serve as explicit cues to accentuate the retrieval of associated pairs and facilitate the associative inference. Xiong et al. (2023) also compared participants' recognition and perceived accuracy rating of associatively inferred fake news with and without hashtags, but did not find significant differences. Because Xiong et al. only evaluated one tweet format, we further investigated the impact of hashtags using different presentation modes.

Considering that relatively few studies have been conducted on people's susceptibility to associatively inferred fake news and there are inconclusive results on the false belief measure and the impact of hashtags, we address three research questions (**RQs**):

- **RQ1:** Given a piece of fake news, do participants recognize it more and give higher accuracy rating when it is associatively inferred from real news than otherwise (e.g., association only or irrelevant)?
- **RQ2:** Do participants' cognitive ability levels impact the effect of associative inference on their susceptibility to misinformation?
- **RQ3:** Do hashtags' presentation modes have an influence on the effect of associative inference on participants' susceptibility to misinformation?

We conducted three online human-subject experiments. Each experiment utilized a two-phase setup. In Experiment 1, we examined whether people's recognition rate and perceived accuracy rating of fake news (Phase 2) depend on how the fake news is associated with real news that they viewed previously (Phase 1) across three conditions: asso-

ciative inference ($a.Inf$), association only ($a.Only$), and control (CON). Between the two phases, participants' cognitive ability level was evaluated to determine whether it moderates the effects of associative inference. We evaluated the effect of embedded hashtags in Experiment 2 using the same setting as Experiment 1. Experiment 3 was designed to replicate the results of Experiments 1 and 2.

In Experiments 1 and 3a, participants in the $a.Inf$ condition, regardless of their cognitive ability levels, showed higher recognition rate and gave higher accuracy rating for the same fake news at Phase 2 than participants in the other conditions (**RQ1 & RQ2**). However, when keywords were presented as hashtags and embedded within tweets (Experiments 2 and 3b), the effect of associative inference was only evident for participants of higher cognitive ability (**RQ3**). Based on our findings, we end the paper by discussing theoretical and practical implications for understanding and mitigating misinformation on social media platforms.

Related Work

Misinformation has been defined as an umbrella term to include any information spreading on social media that is false or inaccurate (Wu et al. 2019). Closely related, Lazer et al. (2018) defined fake news as false or fabricated information written and published to mimic legitimate news media content in form. Given the similar, broad definitions of misinformation and fake news, we use those two terms interchangeably (Lewandowsky and Van Der Linden 2021).

Associative Inference

The associative inference is an adaptive process that allows people to link together related information acquired to make novel connections that they have not directly experienced (Zeithamova and Preston 2010). For example, Carpenter and Schacter (2017) conducted four experiments showing that if participants learned direct associations between two items (AB , e.g., a person [A] and a toy [B] in a room) and then learned direct associations that include one member of the previously studied pairs (BC , e.g., the toy [B] with a different person [C] in a room). The participants were susceptible to misattribute AB event with BC event or vice versa, suggesting the impact of associative inference (i.e., AC).

Lee et al. (2020) conducted a preliminary online study investigating the effect of associative inference on individuals' susceptibility to fake news. Across two phases, they examined participants' recognition and perceived accuracy of news (Phase 2) as a function of how those pieces of news are associated with real news that participants viewed before (Phase 1). The association was varied in three between-subjects conditions (i.e., associative inference, with association, and control) across two phases. Participants' cognitive ability was also measured. Lee et al. only obtained non-significant results, which showed that participants, especially those of higher cognitive ability, tended to give higher perceived accuracy ratings to fake news with associative inference than that without associative inference.

Using the same two-phase procedure, Xiong et al. (2023)

conducted two online human-subjects experiments examining the effect of associative inference with a within-subjects design. Critically, each participant viewed three types of news articles (i.e., real, fake, and fake with associative inference) at Phase 2. In both experiments, the participants recognized more the fake news with associative inference than that without, indicating their susceptibility to associatively inferred fake news. However, the participants gave the lowest accuracy rating for fake news with associative inference than the other types. The authors further analyzed factors that the participants considered for the accuracy ratings of the fake news with associative inference. They proposed that one possible reason could be the *within-subjects* design: Each participant saw fake news with and without associative inference, thus her or his awareness of the associative inference was increased. While the increased awareness could have enhanced the relatively intentional control and effortful process of accuracy rating task (e.g., semantic judgment), it had limited impacts on the relatively intuitive and automatic recognition task (e.g., familiarity judgment).

Cognitive Ability

Prior studies have evaluated participant’s critical thinking ability with the Cognitive Reflection Test (CRT) (Frederick 2005). Participants who scored higher on CRT were also better at discerning fake and real news (Bago, Rand, and Pennycook 2020). The Wordsum has been used to assess people’s cognitive ability in various fields of social science (Malhotra, Krosnick, and Haertel 2007). Recently in the fake news study, Murphy et al. (2019) evaluated participants’ cognitive ability using the Wordsum (Wechsler 2008) and obtained similar findings as using the CRT.

Compared to the CRT, substantially fewer MTurk workers have been previously exposed to the CRT-2 questions (Thomson and Oppenheimer 2016). Lee et al. (2020) and Xiong et al. (2023) evaluated participants’ cognitive ability with the CRT-2 and the Wordsum. Both studies revealed the impact of participants’ cognitive ability on the perceived accuracy ratings, but in the opposite direction. While participants of higher cognitive ability gave lower perceived accuracy ratings for associatively inferred fake news in Xiong et al., Lee et al. only obtained a nonsignificant trend that participants of higher cognitive ability in the associative inference condition tended to give higher perceived accuracy ratings.

Effect of Hashtags

Hashtags are user-created keywords starting with the prefix pound symbol, #, to annotate, categorize and contextualize posts on social media platforms such as Twitter (Huang, Thornton, and Efthimiadis 2010). Previous studies have investigated the impact of hashtags on people’s susceptibility of misinformation on social media but revealed mixed results. Rho and Mazmanian (2019) examined the impact of political hashtags in posts on Meta via a human-subject online study. They found that participants who saw political posts with hashtags were less motivated to know more about related issues compared to participants who viewed the same posts without hashtags. On the contrary, Xiong et al. (2023)

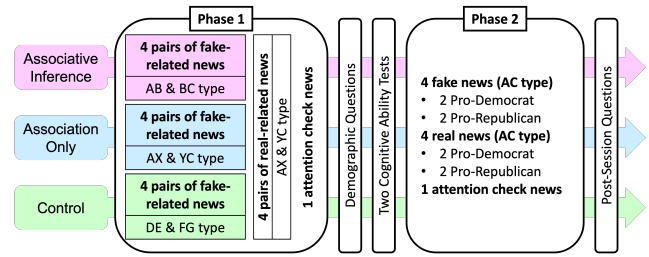


Figure 2: An overview of flow chart for all experiments. “Phase 1” and “Phase 2” boxes show different types of news at each phase. Fake-related, real news at Phase 1 are different among the three conditions (associative inference, association only, and control). Real-related, real news at Phase 1 and all news at Phase 2 are the same across the conditions. Participants answered demographic questions and completed cognitive ability tests between the two phases. We also asked post-session questions at the end.

presented keywords of the snippets of real and fake news as hashtags in tweets. The authors compared participants’ recognition rates and accuracy ratings of fake news with hashtags and without hashtags conditions, but did not find any significant differences. While Rho and Mazmanian embedded bold hashtags in the title of news posts, Xiong et al. only presented the gray hashtags at the bottom of the tweets.

Present Study

In the following, we present three online human-subject experiments investigating whether participants’ recognition rate and perceived accuracy rating of fake news (Phase 2) depend on how the fake news is associated with real news that they viewed previously (Phase 1). In each experiment, there were three between-subjects conditions: associative inference, association only, and control (see Figure 2). Experiment 1 used a tweet format in which gray hashtags were placed at the bottom (see Figure 1). With the effect of associative inference evident in both recognition and perceived accuracy measures in Experiment 1, Experiment 2 applied another format in which blue-highlighted hashtags were embedded in tweet messages. Surprisingly, we did not obtain the effect of associative inference for either measure. Experiment 3 provided a further investigation of the impact of hashtag presentation on associative inference. We replicated the results of Experiments 1 and 2 in Experiments 3a and 3b, respectively.

Experiment 1

Experiment 1 was conducted to investigate **RQ1** and **RQ2**. We anticipated that participants in the associative inference condition would recognize more and give higher accuracy ratings for fake news than those in the other conditions. Awareness of associative inference was expected to be higher for participants with higher cognitive-ability level (Xiong et al. 2023), which should result in the effect of associative inference being more evident for participants with lower cognitive-ability level.

Item	Options	EXP1 (686)	EXP2 (718)	EXP3a (788)	EXP3b (740)
Gender	Male	52.8%	52.0%	51.4%	46.0%
	Female	47.1%	47.6%	48.2%	53.4%
	Other	0.0%	0.3%	0.3%	0.5%
	Prefer NTA	0.1%	0.1%	0.1%	0.1%
Age	18-27	17.8%	17.3%	19.5%	19.7%
	28-37	42.1%	42.8%	43.0%	40.1%
	38-47	18.7%	22.1%	20.9%	20.8%
	48-57	12.1%	13.1%	9.5%	11.1%
	58 or older	9.3%	4.7%	6.7%	8.2%
	Prefer NTA	0.0%	0.0%	0.3%	0.0%
Education Level	No high school	0.3%	0.3%	0.4%	0.3%
	High school	8.3%	6.1%	8.5%	8.2%
	Bachelor	72.9%	68.8%	68.9%	69.3%
	Advanced	18.5%	24.8%	22.0%	22.2%
	Prefer NTA	0.0%	0.0%	0.3%	0.0%
Political Stance	Liberal	48.3%	38.7%	44.0%	49.5%
	Moderate	24.5%	21.5%	23.9%	24.1%
	Conservative	27.3%	39.8%	32.1%	26.4%
	Prefer NTA	0.0%	0.0%	0.0%	0.0%

Table 1: Participants’ demographic information in each experiment (EXP). Number in the bracket on the top row indicates the number of participants. “Prefer NTA” means “Prefer Not to Answer.”

Method

Participants. For the perceived accuracy rating of fake news, a small effect size of the two-way interaction of *cognitive ability* \times *condition*, $\eta_p^2 = .019$, was reported (Lee et al. 2020). Power analysis using G*Power 3.1 (Faul et al. 2007) suggested $n = 501$ participants to detect a small effect size (Cohen’s $f = 0.14$) of interaction of cognitive ability (*lower, higher*) \times condition (*a.Inf, a.Only, CON*) with a power of 0.8 [analysis of variances (ANOVA) test, $\alpha = .05$].

Considering the relatively less control of online experiments on Amazon Mechanical Turk (MTurk), we recruited 1,200 participants to make the design more powerful. Each MTurk worker were only allowed to participate in one of the experiments. The human intelligent tasks (HITs) were posted with the restrictions to workers who (1) were at least 18 years old, (2) previously completed more than 100 HITs and had a HIT approval rate of at least 95%, and (3) were located in the United States. Participants were excluded if they failed one of the two attention checks (Hauser and Schwarz 2016).¹ About 64.1% of the participants passed all attention check questions and completed the study. The median completion time of the experiment was about 12 min. We paid \$1.50 for participants who completed the whole study in this and the rest experiments.

After removing duplicate responses and participants who selected “Prefer not to answer”,² we included 686 partici-

¹We provided instructions asking participants to choose one specified answer to pass the attention check for both recognition and accuracy rating tasks at each phase.

²The removed duplicate responses were from the same IP addresses. We applied the same criteria for all experiments. Due to space limit, we present the exclusion details of this and the following experiments in Table S1 of the Supplementary Materials.

pants’ results for the data analysis at Phase 2 and 644 participants’ answers were used for the data analysis at Phase 1. Participants’ demographics are shown in Table 1.

Materials. A total of 42 different news were presented. Each news was a snippet from one real or fake news article. Thirty-eight pieces of the news were based on real news articles reported from major news media, including *washingtonpost.com*, *usatoday.com*, and *foxnews.com*. The remaining four pieces of news were based on the fake news that were debunked by the fact-checking website *snopes.com*.

Each news was designated by two letters (e.g., *AB, YC* or *DE*). Each letter, such as *A, Y*, and *D*, refers to a unique keyword (e.g., one public figure, an event, or an entity, see Figure 1 as an example) in the news. Each keyword was also listed after the pound symbols as hashtags in the news. To minimize the impact of source (Visentin, Pizzi, and Pichierri 2019), we applied an identical Twitter user ID and a blurred profile image for each piece of news. Stimuli in all experiments are shown in the Supplementary Figures S1-S24.³

Due to the main interest in associative inference, we started the news selection from the four pieces of fake news (i.e., *AC*) at Phase 2. Across the three conditions, the fake-news triads between the two phases are as follows:

- *a.Inf* ($AB \& BC \rightarrow AC$): the four pairs of fake-related, real news at Phase 1 were in an *AB & BC* type such that two keywords of each news overlapped through one common keyword *B* (e.g., *AB*: Mitch McConnell [*A*], Ping May [*B*]; *BC*: Ping May [*B*], cocaine [*C*]), affording an associative inference for the fake news at Phase 2 (e.g., Mitch McConnell [*A*], cocaine [*C*], see Figure 1).
- *a.Only* ($AX \& YC \rightarrow AC$): the four pairs of fake-related, real news at Phase 1 were in an *AX & YC* type, each of which had no overlapped keyword (e.g., Mitch McConnell [*A*], Climate Change [*X*], Authorities [*Y*], and Cocaine [*C*]). Thus, each piece of real news was only associated with the fake news at Phase 2.
- *CON* ($DE \& FG \rightarrow AC$): the four pairs of fake-related, real news at Phase 1 were in a *DE & FG* type. Each pair have neither an overlapped keyword nor was associated with the fake news at Phase 2.

The four real-news triads were constructed in the same way as the fake-news triads in the *a.Only* condition. Since news at Phase 2 was politically related, the news stance was also controlled at each veracity level: half was pro-Republican and the other was pro-Democrat. Two extra pieces of real news, one for each phase, were used for attention check.

Procedure. After participants accepted the HIT on MTurk, they were directed to the online survey designed on Qualtrics. Participants who agreed with the consent form were randomly assigned to one of three conditions (*a.Inf, a.Only, CON*), the procedure of which was identical.

Phase 1. Eight pairs of real news were presented in a randomized order in each condition. Half of the pairs were related to real news at Phase 2 (i.e., real-related) and the

³All Supplementary Materials for this paper can be accessed through the following link: <http://tiny.cc/Leeetal2023supp>

other half were related to fake news at Phase 2 (i.e., fake-related). We also counterbalanced the presentation order of two pieces of news in each pair.

After participants viewed each piece of news, they were asked to answer, “Have you ever seen this before?” with four options (*Yes, Unsure, No, and Prefer not to answer*) at first. Such recognition task was used to reflect whether participants have been exposed to the news before our study. Then participants answered, “How would you rate the accuracy of this news article?” using a 5-point scale (1: *Very inaccurate*, 5: *Very accurate*). For the perceived accuracy ratings, we also provided “*Prefer not to answer*” as the 6th option. The perceived accuracy rating thus evaluated whether participants were able to identify fake news.

Between Two Phases. After Phase 1, participants filled out their demographic information, such as age, gender, and educational background. We also assessed participants’ cognitive ability using the CRT-2 (Thomson and Oppenheimer 2016) and the Wordsum (Wechsler 2008). CRT-2 measured participants’ tendency to override an incorrect “gut” response with four questions. For example, we asked participants, “If you are running a race and you pass the person in the second place, what place are you in?” While the intuitive answer is “*first*”, the correct answer is “*second*.” The Wordsum measured participants’ intelligence scale of vocabulary with ten items. Within the test, we showed participants different words in capital letters (e.g., SPACE). Then, we asked participants to choose one word that comes closest to the meaning of the word in capital letters from five options (e.g., captain, school, noon, board, room). Options of “*do not know*” and “*Prefer not to answer*” were also provided for both tests.

Phase 2. After the demographic questionnaire and cognitive ability tests, Phase 2 started, in which we presented eight pieces of news (half fake and half real) in a randomized order. Same as Phase 1, participants completed the recognition and perceived accuracy tasks for each piece of news.

Post-session Questionnaire. After Phase 2, we asked participants’ interests in politics and their political ideology. Additionally, we randomly presented one piece of the four fake news at Phase 2 with the perceived accuracy rating that the participant gave and asked the participant to select all factors affecting his/her decision. Seven options (*Source, Writing style, Content, Web search results*⁴, *News presented in Stage 1, News that I saw before this study, and Opinions from others*) were provided which were presented in a randomized order. We also provided “*Other (with text box)*” and “*Prefer not to answer*” as two extra options in the end. To further understand the effect of associative inference, for any participant who chose “*News presented in Stage 1*” (Stage 1 refers to Phase 1 in the paper), we asked the participant to explain how the “*News presented in Stage 1*” had influence on his/her decision with an open-ended question.

⁴We could not prevent participants from using web search in the online survey, so we included “web search results” as an option. There were 6.85% (47) of participants selected this option.

Data Analysis

We used the same method to analyze the recognition and perceived accuracy results at Phases 1 and 2 for this and the following experiments. All independent variables and dependent measures in our experiments are described in Supplementary Table S4. In literature, participants 1) were categorized as of high cognitive ability with eight or more items answered correctly for the Wordsum (Murphy et al. 2019) and 2) got 56.2% of the four items answered correctly for the CRT-2 on average (Thomson and Oppenheimer 2016). Thus, we chose *ten* as a cutoff to categorize participants of a higher or lower cognitive-ability-test score for a total of 14 questions of the Wordsum and the CRT-2. We categorized 419 participants who gave ten or more correct answers as having a *higher* cognitive-ability-test score and the remaining 267 participants as having a *lower* score. There were 231, 217, and 238 participants in the *a.Inf*, *a.Only*, and *CON* conditions, respectively. The demographics were similar across conditions (see Table S2 in the Supplementary Materials).

Phase 1. Each pair of real news was related to one piece of real or fake news in Phase 2. Thus, we divided news in Phase 1 into two types, *real-related* and *fake-related*. Recognition rates and perceived accuracy ratings were entered into 2 (*news type: real-related, fake-related*) \times 3 (*condition: a.Inf, a.Only, CON*) \times 2 (*cognitive-ability-test score: lower, higher*) mixed ANOVAs with a significance level of .05, respectively. Post-hoc tests with Bonferroni correction were performed, testing all pairwise comparisons with corrected *p* values for possible inflation. We report the effect size using η_p^2 (Keppel 1991), which is reported by SPSS.

Phase 2. Recognition rates and perceived accuracy ratings were entered into 2 (*news veracity: real, fake*) \times 3 (*condition: a.Inf, a.Only, CON*) \times 2 (*cognitive-ability-test score: lower, higher*) mixed ANOVAs with a significance level of .05. Post-hoc analyses were conducted in the similar way to the results of Phase 1.

We further performed Linear Mixed Effect Regression (LMER) on perceived accuracy ratings at Phase 2 for each experiment with the lme4 package in R (Bates et al. 2011). LMER allows controlling random effects of participants and news trials (Brauer and Curtin 2018). We intentionally used the same model as ANOVAs for LMERS, and included random intercepts for participants and news trials. Thus, we could compare the results of ANOVAs and LMERS. Following ANOVA results, we report LMER results in square bracket without the effect size (Pennycook et al. 2021), because it is under debate for a single, agreed standardized effect size of LMER (Rights and Sterba 2018). We report degrees of freedom with Satterthwaite approximation. ANOVAs and LMERS show similar results. Supplementary Tables S25–S28 show the complete results of LMERS.

Thematic Analysis. We did a thematic analysis (Braun and Clarke 2006) for the answers of open-ended questions. The first two co-authors of this article and two undergraduate students performed the thematic analysis independently at first. Then they discussed the results and finalized the thematic analysis together.

Effect		Experiment 1						Experiment 2							
		df	Recognition			Accuracy			df	Recognition			Accuracy		
			F	p	η_p^2	F	p	η_p^2		F	p	η_p^2	F	p	η_p^2
Phase 1	News type	1, 638	34.18	.000	.051	0.56	.453	.001	1, 647	16.17	.000	.024	1.91	.167	.003
	News type*Cond.	2, 638	4.37	.013	.013	14.27	.000	.043	2, 647	1.95	.143	.006	8.58	.000	.026
	News type*Cog.	1, 638	7.16	.008	.011	1.59	.208	.002	1, 647	5.10	.024	.008	0.07	.792	.000
	News type*Cond.*Cog.	2, 638	0.29	.750	.001	0.95	.388	.003	2, 647	0.06	.945	.000	1.41	.246	.004
	Cond.	2, 638	2.14	.118	.007	1.96	.141	.006	2, 647	0.51	.600	.002	1.81	.164	.006
	Cog.	1, 638	54.01	.000	.078	2.44	.119	.004	1, 647	95.99	.000	.129	17.46	.000	.026
Phase 2	Cond.*Cog.	2, 638	0.23	.797	.001	1.50	.223	.005	2, 647	0.65	.523	.002	0.10	.903	.000
	News vera.	1, 680	5.08	.025	.007	524.71	.000	.436	1, 712	10.28	.001	.014	370.29	.000	.342
	News vera.*Cond.	2, 680	3.35	.036	.010	5.98	.003	.017	2, 712	1.07	.342	.003	1.49	.227	.004
	News vera.*Cog.	1, 680	0.57	.451	.001	43.38	.000	.060	1, 712	2.24	.135	.003	119.87	.000	.144
	News vera.*Cond.*Cog.	2, 680	0.30	.738	.001	0.90	.406	.003	2, 712	0.02	.982	.000	2.12	.121	.006
	Cond.	2, 680	0.36	.698	.001	1.19	.304	.003	2, 712	0.93	.396	.003	2.18	.113	.006
	Cog.	1, 680	34.89	.000	.049	29.14	.000	.041	1, 712	114.67	.000	.139	138.03	.000	.162
Cond.*Cog.	2, 680	0.92	.399	.003	1.00	.369	.003	2, 712	0.25	.780	.001	0.35	.704	.001	

Table 2: Summary of the ANOVA results in Experiments 1 and 2. Note. “df” denotes “degrees of freedom” that is the same for the recognition and accuracy measures in each phase. Bold font indicates statistical significance ($p < .05$). “Cond.” means “Condition,” “Cog.” means “Cognitive-ability-test score,” and “vera.” means “veracity.”

Results

Figure 3 displays the averages of both measures obtained for each condition, and Table 2 presents the results of the ANOVAs. We provide the complete post-hoc analysis results in the Supplementary Materials.

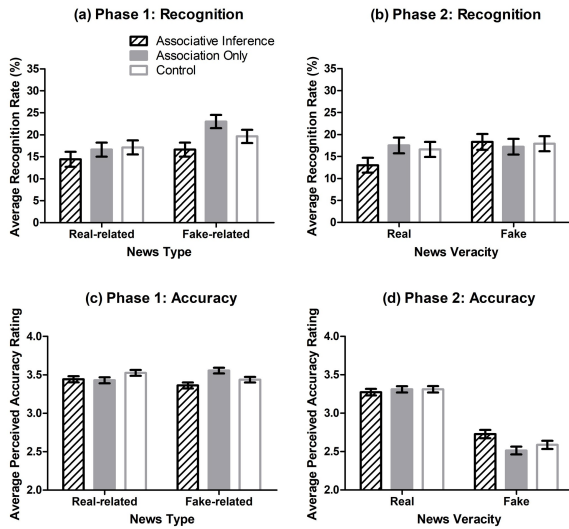


Figure 3: Average recognition rate (top row) and perceived accuracy rating (bottom row) across the three conditions for each news type at Phase 1 (left column) and each news veracity at Phase 2 (right column) of Experiment 1. Error bars represent \pm one standard error.

Phase 1. Figures 3a and 3c depict the results of Phase 1, which serve as baseline to understand possible differences across the conditions *before* our key manipulation. Interestingly, participants recognized more fake-related news (19.7%) than real-related news (16.0%). The significant main effect of news type (see Table 2) suggests that fake news might have been based on real news that is more familiar to news consumers. Such a higher recognition rate of the fake-related news was more evident for participants

of a higher cognitive-ability-test score (fake-related vs. real-related: 14.3% vs. 8.9%) than those of a lower score (fake-related vs. real-related: 25.2% vs. 23.1%).

The two-way interaction of *news type* (2: real-related, fake-related) \times *condition* (3: *a.Inf*, *a.Only*, *CON*) was significant for both recognition rates and the perceived accuracy ratings (see Table 2). For the real-related news, neither recognition rate nor perceived accuracy measure showed any statistically significant difference across the conditions (see Figures 3a and 3c). In contrast, for the fake-related news, the differences across conditions were evident for both recognition rates and perceived accuracy ratings. Specifically, participants in the *a.Only* condition recognized more the fake-related news than those in the *a.Inf* condition ($p_{adj} = .004$), and gave higher accuracy rating in general ($p_{adj} \leq .023$). Such results are essential, which indicate that any higher susceptibility to associatively inferred fake news at Phase 2 is not due to participants’ higher familiarity or more belief in fake-related real news at Phase 1.

Phase 2. We presented the same fake and real news to the participants in each condition. Thus, any differences across the conditions would indicate the impacts of key manipulations across Phases 1 and 2.

Participants recognized more fake news (17.8%) than real news (15.7%), but such a main effect of *news veracity* was *condition* dependent (see Figure 3b and Table 2). Post-hoc analysis revealed that only participants in the *a.Inf* condition recognized more the fake news (18.3%) than the real news (13.0%, $F_{(1,229)} = 9.36, p = .002, \eta_p^2 = .039$). Thus, the effect of associative inference was evident for the familiarity judgment, which is consistent with Xiong et al. (2023).

Turning to the accuracy measure (see Figure 3d), participants gave higher ratings for the real news (3.30) than for the fake news (2.61), indicating that they could differentiate the fake news from the real news. Same as the recognition measure, the two-way interaction of *news veracity* \times *condition* was also significant ($F_{(2,680)} = 5.98, p = .003, \eta_p^2 = .017$, [$F_{(2,4790)} = 7.17, p = .001$]). Across the three conditions, participants’ accuracy ratings were similar for the real news,

but showed differences for the fake news (*a.Inf* vs. *a.Only* vs. *CON*: 2.73 vs. 2.51 vs. 2.59, $F_{(2,680)} = 4.08, p = .017, \eta_p^2 = .012$). Post-hoc analysis revealed that, for the same fake news, the average accuracy rating in the *a.Inf* condition was significantly higher than that of the *a.Only* condition ($p_{adj} = .005$), and showed a non-significant trend to be higher than that of the *CON* condition ($p_{adj} = .059$).

Finding 1: *Using a between-subjects design, we found the effect of associative inference for both recognition and perceived accuracy measures (RQ1).*

The main effect of cognitive-ability-test score was significant for both recognition rates and perceived accuracy ratings (see Table 2). Participants of a lower cognitive-ability-test score recognized more news (22.1%) and gave higher accuracy ratings (3.08) than those of a higher cognitive-ability-test score (recognition: 11.5%, accuracy: 2.83). The two-way interaction of *news veracity* \times *cognitive-ability-test score* was also significant for the perceived accuracy measure. Post-hoc analysis revealed that participants gave similar accuracy ratings for the real news regardless of their cognitive-ability levels. In contrast, for the fake news, participants of a higher cognitive-ability-test score gave lower accuracy rating (2.39) than participants of a lower cognitive-ability-test score (2.83, $F_{(1,680)} = 53.39, p < .001, \eta_p^2 = .073$). Thus, participants with higher cognitive-ability level were less susceptible to misinformation than those with lower cognitive-ability level in general (Pennycook and Rand 2019). Yet participants' cognitive-ability level showed no impact on the obtained effects of associative inference.

Thematic Analysis. Among the 686 participants, 76 (11.1%) of them indicated that their accuracy ratings were impacted by "News presented in Stage 1." We further disregarded 24 meaningless answers for the open-ended question, such as "yes" or "nice." For the remaining 52 meaningful answers, most of them were in the *a.Inf* condition (73.1%), but only 11.5% and 15.4% were in the *a.Only* and *CON* conditions, respectively. Such results highlight that more participants in the *a.Inf* condition explicitly made their accuracy decision of the fake news based on "News presented in Stage 1" than those in the other conditions. Supplementary Table S3 shows the percentage of each factor affecting participants' perceived accuracy rating decision.

We identified three major themes: 1) *Association between Two Phases*: Among the 52 answers, 46.2% (24) of the participants indicated that when they gave accuracy ratings for fake news in Phase 2, they noticed its connection with news headlines in Phase 1. For example, one participant answered, "Because it made me think that if it were repeated more than once there was some possibility it was accurate." 2) *Gist-based Recall*: 28.8% (15) of the participants made the accuracy rating decision based on their prior knowledge of or belief in the news. For instance, one participant explained, "I remember seeing it on a Facebook post, and this is something that I believe can be very true." 3) *Verbatim Recall*: About 23.1% (12) participants noticed the gap between the news in two phases. They recalled the details of news in Phase 1 and detected the distorted or exaggerated parts in Phase 2. For example, one participant answered, "There

were a few modified stage 1 messages that were reworded or combined for stage 2..."

We also examined participants' cognitive-ability-test score distribution in each theme. For the theme of association between two phases, 67% (16) of the participants were of a higher cognitive-ability-test score, showing a similar distribution that includes all participants. For the verbatim recall theme, 92% (11) of the participants were of a higher cognitive-ability-test score, but the percentage was only 40% (6) for the theme of gist-based recall.

Finding 2: *Individuals' cognitive-ability level showed no impact on the effects of associative inference (RQ2). Responses to the post-session questions revealed that few participants were aware of the association or inference.*

Experiment 2

Hashtags are typically highlighted in blue and embedded within the tweet messages. Thus, we recruited extra participants in Experiment 2 and evaluated whether the effects of associative inference obtained in Experiment 1 (**RQ1** & **RQ2**) can be generalized to or accentuated with blue-highlighted hashtags embedded in tweets (**RQ3**).

Participants, Stimuli, and Procedure

We published 1,200 tasks on MTurk. Applying the same criteria as Experiment 1, we included 718 participants' results for the data analysis in Phase 2. There were 255 participants who were categorized as with a *higher* cognitive-ability-test score and the remaining 463 participants were categorized as with a *lower* cognitive-ability-test score. The median completion time was about 14 min. We analyzed 653 participants' results in Phase 1.

We conducted Experiment 2 using the same setting as Experiment 1 except as noted. In Experiment 2, we presented hashtags in blue and embedded the hashtags within the tweets to assess whether such a presentation would impact the effect of associative inference and participants' susceptibility to fake news. Since few participants were aware of the connection between phases (see Thematic Analysis section of Experiment 1), we decided not to ask the open-ended question in this and the following experiments.

Results

Phase 1. The results are shown in Figures 4a and 4c, which replicate the main findings of Experiment 1. As expected, participants recognized more the fake-related news (27.3%) than the real-related news (24.7%), but such a recognition gap did not show any significant differences across conditions. For the perceived accuracy rating, only the two-way interaction of *news type* \times *condition* was significant (see Table 2). Same as Experiment 1, post-hoc analysis showed that differences across the conditions were only significant for the fake-related news but not for the real-related news. Specifically, participants in the *a.Inf* condition (3.35) gave lower perceived accuracy rating than those in the *CON* condition (3.53, $p_{adj} = .011$). The main effect of cognitive-ability-test score was significant for both measures. Participants of a lower cognitive-ability-test score recognized more

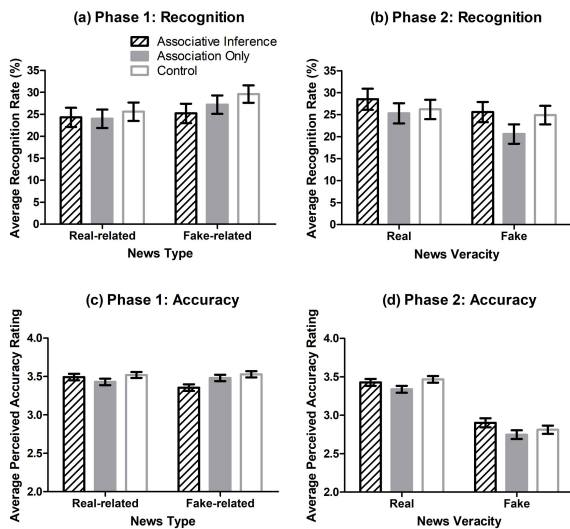


Figure 4: Average recognition rate (top row) and perceived accuracy rating (bottom row) across the three conditions for each news type at Phase 1 (left column) and each news veracity at Phase 2 (right column) of Experiment 2. Error bars represent \pm one standard error.

news (37.6%) and gave higher accuracy ratings (3.56) than those of a higher score (recognition: 14.4%, accuracy: 3.38). The interaction of *cognitive-ability-test score* \times *news type* was only significant for the recognition measure. Specifically, a non-significant trend of higher recognition rate of fake-related news than real-related news was more evident for participants of a higher cognitive-ability-test score (fake-related vs. real-related: 16.4% vs. 12.3%) than those of a lower score (fake-related vs. real-related: 38.2% vs. 37.0%).

Phase 2. As illustrated in Figures 4b and 4d, the main effect of *news veracity* was significant for both recognition and perceived accuracy measures. Same as Experiment 1, participants could differentiate the real news (3.41) from the fake news (2.82). In contrast, they recognized more the real news (26.7%) than the fake news (23.7%) in Experiment 2. The main effect of *cognitive-ability-test score* was also significant for both measures. Participants of a lower cognitive-ability-test score recognized more (38.2%) and gave higher accuracy ratings (3.41) than those of a higher score (recognition: 12.2%, accuracy: 2.82). Also, the interaction of *news veracity* \times *cognitive-ability-test score* was significant for the accuracy rating (see Table 2). Specifically, participants of a higher cognitive-ability-test score could distinguish the real news (3.29) from the fake news (2.36), but those of a lower score rated news as real in general (real: 3.54, fake: 3.28).

Between-experiments Comparison. Different from Experiment 1, *condition* did not show any significant impacts on either measure. One possible explanation is offered by the encoding of associated pairs in human memory (Craik et al. 1996). Highlighting relevant parts in the texts has often been considered an effective encoding process tool to aid memory (Lorch 1989). Yet preexisting highlighting could interfere with reading comprehension, resulting in decreased performance (Hunt and Lamb 2001). Peterson (1992) exam-

ined college students' learning of a history chapter. Results showed that students who highlighted while reading performed worse on the tests of comprehension wherein they needed to make inferences that required connecting different ideas across the text. Prior works have supported the notion that the effectiveness of encoding can be determined by the presented materials (Brown and Craik 2000). Participants in Experiment 2 might have primarily attended to the blue-highlighted hashtags in tweets, but encoded the rest parts of the tweets shallowly, thus impeding associative inference.

Finding 3: *After embedding blue-highlighted hashtags within tweet messages, we obtained the results of Phase 1 similar to those of Experiment 1. Yet the non-significant interactions with condition at Phase 2 indicate that the blue-highlighted hashtags embedded within tweet messages might have reduced the associative inference (RQ3).*

Experiment 3

We conducted two sub-experiments to further understand the impacts of hashtags presentation on associative inference (RQ3). Experiment 3a was identical to Experiment 1 except that the hashtags at the bottom of the tweets were highlighted in blue (see Figure 5a). We evaluated the impact of blue-highlighted hashtags on the effect of associative inference. Experiment 3b was the same as Experiment 2 except that the hashtags were presented at the bottom of the tweets again (see Figure 5b). We tested whether presenting hashtags again was critical to obtain the effect.

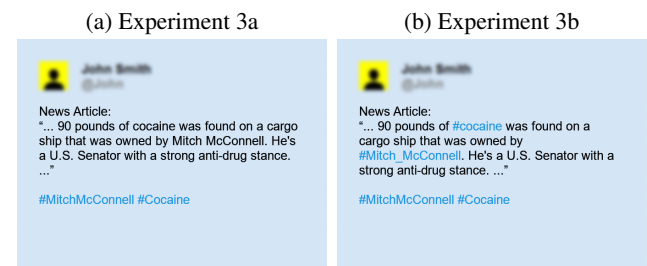


Figure 5: An example of news stimulus in Experiments 3a (left panel) and 3b (right panel).

Participants, Stimuli, and Procedure

We published 2,400 HITs on MTurk. We included 1,528 participants' results for the data analysis at Phase 2. Among them, 51.6% (788) were in Experiment 3a and 48.4% (740) were in Experiment 3b. The demographics of each experiment are shown in Table 1. We analyzed 1,427 participants' results of Phase 1, with 51.3% (732) in Experiment 3a and 48.7% (695) in Experiment 3b. The median completion time was about 13 min for both experiments.

Results of Experiments 3a & 3b

We analyzed the results of Experiments 3a and 3b using the same methods as the prior experiments. In Experiment 3a, 423 and 365 participants were categorized as having a *higher* and *lower* cognitive-ability-test score, respectively.

Effect		Experiment 3a							Experiment 3b						
		df	Recognition			Accuracy			df	Recognition			Accuracy		
			F	p	η_p^2	F	p	η_p^2		F	p	η_p^2	F	p	η_p^2
Phase 1	News type	1, 726	51.37	.000	.066	0.34	.560	.000	1, 689	12.44	.000	.018	15.87	.000	.023
	News type*Cond.	2, 726	2.91	.055	.008	13.71	.000	.036	2, 689	3.29	.038	.009	12.42	.000	.035
	News type*Cog.	1, 726	2.13	.145	.003	1.14	.286	.002	1, 689	15.34	.000	.022	0.11	.738	.000
	News type*Cond.*Cog.	2, 726	0.94	.392	.003	0.63	.534	.002	2, 689	0.30	.740	.001	0.34	.709	.001
	Cond.	2, 726	0.24	.784	.001	0.35	.702	.001	2, 689	0.04	.960	.000	0.27	.765	.001
	Cog.	1, 726	59.46	.000	.076	2.24	.135	.003	1, 689	76.80	.000	.100	0.07	.786	.000
Cond.*Cog.	2, 726	3.30	.037	.009	1.13	.325	.003	2, 689	0.78	.459	.002	4.75	.009	.014	
Phase 2	News vera.	1, 782	3.61	.058	.005	698.48	.000	.472	1, 734	10.98	.001	.015	624.81	.000	.460
	News vera.*Cond.	2, 782	8.08	.000	.020	7.69	.000	.019	2, 734	1.11	.330	.003	0.47	.623	.001
	News vera.*Cog.	1, 782	0.92	.338	.001	35.56	.000	.043	1, 734	2.12	.146	.003	55.51	.000	.070
	News vera.*Cond.*Cog.	2, 782	0.61	.544	.002	0.57	.564	.001	2, 734	3.23	.040	.009	2.43	.089	.007
	Cond.	2, 782	0.89	.409	.002	6.65	.001	.017	2, 734	1.40	.248	.004	4.02	.018	.011
	Cog.	1, 782	74.60	.000	.087	40.67	.000	.049	1, 734	68.15	.000	.085	56.47	.000	.071
	Cond.*Cog.	2, 782	2.57	.077	.007	1.03	.356	.003	2, 734	1.42	.243	.004	1.02	.361	.003

Table 3: Summary of the ANOVA results in Experiments 3a and 3b. Note. “df” denotes “degrees of freedom” that is the same for the recognition and accuracy measures in each phase. Bold font indicates statistical significance ($p < .05$). “Cond.” means “Condition,” “Cog.” means “Cognitive-ability-test score,” and “vera.” means “veracity.”

Similarly, 434 and 306 participants were categorized as having a *higher* and a *lower* cognitive-ability-test score in Experiment 3b. Figure 6 displays the averages of both measures obtained for each experiment. Table 3 shows the results of ANOVAs. We also provide the complete post-hoc analysis in the Supplementary Materials.

Experiment 3a. Experiment 3a replicated all the major findings of Experiment 1. Participants recognized more the fake-related news (20.1%) than the real-related news (16.2%) at **Phase 1** (see Figure 6a). The effect of *news type* showed a non-significant trend to be different across *conditions*. The main effect of *news type* was not significant for the accuracy measure, but its interaction with *condition* was significant (see Table 3). While participants in the *a.Inf* condition perceived higher accuracy for the real-related news than the fake-related news, those in the other conditions revealed an opposite pattern (see Figure 6c).

When moving to **Phase 2** (see Figures 6b, 6d, and Table 3), the two-way interaction of *news veracity* \times *condition* was significant for both recognition ($F_{(2,782)} = 8.08, p < .001, \eta_p^2 = .020$) and accuracy ($F_{(2,782)} = 7.69, p < .001, \eta_p^2 = .019, [F_{(2,5503)} = 10.37, p < .001]$) measures. Results of post-hoc analysis showed that *condition* was only significant for the fake news recognition ($F_{(2,782)} = 3.63, p = .027, \eta_p^2 = .009$) and accuracy rating ($F_{(2,782)} = 10.24, p < .001, \eta_p^2 = .026$). Specifically, participants in the *a.Inf* condition recognized more fake news (20.2%) than those in the *CON* condition (14.1%, $p_{adj} = .030$). Similarly, across the three conditions, participants in the *a.Inf* condition gave the highest accuracy rating (2.80, $p_{adj} \leq .011$).

Experiment 3b. Experiment 3b replicated the results of Experiment 2 (see Table 3). At **Phase 1**, the main effects of *news type* and *cognitive-ability-test score*, as well as the two-way interaction of *news type* \times *cognitive-ability-test score* were significant for the recognition measure. The interaction of *news type* \times *condition* was also significant for the accuracy measure. At **Phase 2**, there were main effects of *news veracity* and *cognitive-ability-test score* for both measures, but the interaction of *news veracity* \times *cognitive-ability-test*

score was only significant for the accuracy measure.

Also, at **Phase 1**, the interaction of *news type* \times *condition* was significant for the recognition measure. Participants recognized more the fake-related news than the real-related news in the *a.Only* condition and *CON* condition, but not in the *a.Inf* condition ($F < 1$, see Figure 6e). Furthermore, participants gave slightly higher accuracy rating for the real-related news (3.47) than the fake-related news (3.40).

At **Phase 2**, in addition to the replicated results, the three-way interaction of *news veracity* \times *condition* \times *cognitive-ability-test score* was significant ($F_{(2,734)} = 3.23, p = .040, \eta_p^2 = .009$) for the recognition measure (see Figures 6f, 6g, and Table 3). Post-hoc analyses revealed that the two-way interaction of *news veracity* \times *condition* was only significant for the participants of a higher cognitive-ability-test score ($F_{(2,431)} = 5.06, p = .007, \eta_p^2 = .023$). Specifically, participants of a higher cognitive-ability-test score in the *a.Inf* condition revealed a non-significant trend of larger recognition rate for the fake news (10.0%) than for the real news (7.4%), but those in *CON* (fake vs. real = 8.3% vs. 11.4%) and *a.Only* (fake vs. real = 6.6% vs. 10.3%) conditions revealed an opposite pattern. For the participants of a lower cognitive-ability-test score, the two-way interaction of *condition* \times *news veracity* was not significant, $F < 1$.

The effect of *condition* was also significant for the perceived accuracy measure (see Figures 6i and 6j). Specifically, participants in the *a.Inf* condition (3.08) gave higher accuracy rating than those in *a.Only* condition (2.93, $p_{adj} = .019$). Also, there was a non-significant three-way interaction of *news veracity* \times *condition* \times *cognitive-ability-test score* ($F_{(2,734)} = 2.43, p = .089, \eta_p^2 = .007, [F_{(2,5168)} = 3.28, p = .038]$). The significant interaction in LMER suggests that the random effect of news trials might be essential to understand the individual differences in such analysis.

Between-experiments Comparison. To check the effect of blue-highlighted hashtags (Exp.1 vs. Exp.3a) and repeated hashtags (Exp.2 vs. Exp.3b), we conducted additional ANOVAs by adding *experiment* as another between-subjects factor for the combined data of Experiments 1 and 3a, and

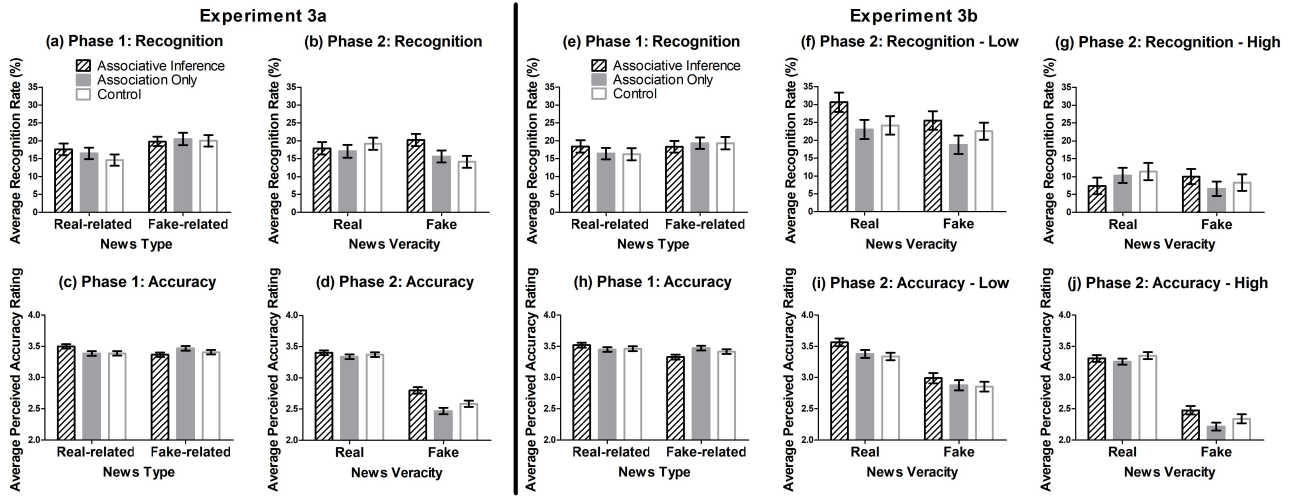


Figure 6: Average recognition rate (top row) and perceived accuracy rating (bottom row) of Experiments 3a and 3b. Phase 2 results of Experiment 3b for participants of a lower cognitive-ability-test score (“Low”) and a higher score (“High”) are presented. Error bars represent \pm one standard error.

Experiments 2 and 3b, respectively. Here, we selectively report the three-way interactions involving *experiment*. Complete analysis results are listed in the Supplementary Materials. The interaction of *experiment* \times *news veracity* \times *condition* was not significant for either recognition or accuracy measures ($F_s \leq 1.68$), indicating limited impacts of blue-highlighted or repeated hashtags when they were presented at the bottom of tweets.

Finding 4: *Experiments 3a and 3b replicated the results of Experiments 1 and 2, respectively. Regardless of experiments, results of Phase 1 kept revealing an appropriate baseline for investigating associative inference at Phase 2. Moreover, participants of higher cognitive ability level might have been less impacted by the presentation modes (e.g., embedded hashtags) and still showed the susceptibility to associatively inferred misinformation (RQ3).*

General Discussion

Findings from our three experiments consistently showed that associative inference could be one cognitive driver of misinformation susceptibility. To this end, we argue that adaptive memory processes (e.g., associative inference) can become maladaptive (e.g., increasing individuals’ familiarity and belief in fake news), highlighting the essentiality of bringing insights from cognitive psychology into the understanding and mitigation of misinformation. We discuss the theoretical and practical implications of the study.

Associative Inference Increases Recognition and Perceived Accuracy Rating. The most clear-cut findings are that associative inference increases individuals’ susceptibility in recognition (i.e., familiarity) and perceived accuracy measures. Thus, our study provides novel evidence showing increased *belief* in associatively inferred fake news, which was not obtained in prior studies (Xiong et al. 2023). Xiong et al. explained that their within-subjects design might have made participants, especially those of higher cognitive abil-

ity level, more aware of the associative inference across phases. Our findings corroborate their account. Using a between-subjects design, we found that the increased belief in associatively inferred fake news was not dependent on participants’ cognitive-ability levels. Thus, without relative comparison with other conditions (e.g., association only), individuals, even those of higher cognitive ability level, could miss the associative inference across phases.

The effect sizes of associative inference may indicate a small to moderate effect, but the observed magnitude is generally in line with the average effect size in the context of misinformation (Lutzke et al. 2019; Murphy et al. 2019). Small effects can be still worth noting (Lakens 2013), especially considering the issue of misinformation spread out on social media that can be aggregated across individuals (Gelman 2018) and over time (Funder and Ozer 2019).

Impact of Hashtags’ Presentation. Comparing to the color coding, embedding hashtags in the tweet messages turned out to be more influential for the effect of associative inference. Highlighted hashtags in a tweet message could draw a reader’s attention. Because the hashtags were highlighted, participants of higher cognitive ability might seek to understand why, which required thinking about the meaning of the overall tweet and how its different pieces relate to one another, indicating analytical thinking (Worthen et al. 2006). In contrast, participants of lower cognitive ability might pay attention to the “pop-out” hashtags mainly, and were too “lazy” to go through the meaning of the overall tweets. Therefore, individuals who are better equipped with adaptive ability to combine existing knowledge in response to novel circumstances can be more vulnerable to misinformation. This is a post-hoc explanation, and future research needs to more thoroughly replicate and investigate the difference.

Associatively Inferred Misinformation Mitigation. The difference between the two designs (i.e., within-subjects and between-subjects) is particularly important for generating

guidelines to mitigate associatively inferred misinformation. Existing misinformation correction often focuses on correcting inaccurate contents and providing accurate information (Barrera et al. 2020) or information source (Bode and Vraga 2018). Yet the gap between our study and Xiong et al. points out an approach of *process-based* correction. For instance, a correction that highlights the association and possible inference between existing real news and misinformation might help online news consumers become aware of the attempt at deception. Such process-based correction might offer protection against different types of misinformation that leverage associative inference.

Empirical Analysis of Associatively Inferred Misinformation. Moreover, developing effective misinformation mitigation requires understanding of how associatively inferred fake news occurred *in practice*. So far, relatively few experimental studies (Lee et al. 2020; Xiong et al. 2023) have examined political misinformation using a specific type (i.e., $AB \& BC \rightarrow BC$), which may not be representative.

Building upon predicting susceptibility to misinformation among social media users (Shen et al. 2019; Teng et al. 2022), future research can conduct an empirical analysis by systematically measuring afforded associative inferences between misinformation on social media and real news consumed by social media users to predict susceptible users. By predicting more vulnerable group of users, researchers and practitioners can set a priority to mitigate the spread of misinformation on social media.

Limitations. Our study has a few limitations. First, all tweets used in our study are mainly related to politics. Further investigation on other topics, e.g., health information, can improve the external validity of our study. Second, we recruited MTurk workers who tended to be younger, better educated, and have better digital literacy (Guess and Munger 2022). Thus, our results may represent a population having more concerns on fake news than the broader U.S. public. Third, our study only examined the impact of associative inference in the short term. Future work on extending the gap between two phases can help reveal whether the effect of associative inference will hold in the long term.

Conclusion

Our study shows that one basis for people's susceptibility to misinformation is adaptive, constructive processes of memory. Fake news can be designed to appeal cognitive preference, maybe more than real news. While individuals of lower cognitive-ability level are more susceptible to misinformation in general, our results indicate that those of higher cognitive-ability level are more susceptible to associatively inferred misinformation. Different from the idea that misinformation is attractive to "lazy" people, our study implies that misinformation, at least some of which fits with our cognitive mechanisms, can make more adaptive individuals also become susceptible.

Ethical Statement

We have taken careful steps to ensure research ethics. First, we worked with our institutional review board (IRB) office

and obtained their approval before running the online experiments. Second, implied consent was obtained for each participant before all experiments. Third, the study presented little to no risk compared to those encountered in people's everyday online activities. Moreover, the results from our study can benefit researchers and practitioners to build better tools to mitigate misinformation from human aspects. We believe the benefit outweighs the potential risk.

Acknowledgements

This research was supported in part by Penn State under PSU SSRI Seed Grant and the National Science Foundation under grants 1820609, 1915801, and 2121097.

References

- Allcott, H.; and Gentzkow, M. 2017. Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31(2): 211–236.
- Bago, B.; Rand, D. G.; and Pennycook, G. 2020. Fake news, fast and slow: Deliberation reduces belief in false (but not true) news headlines. *Journal of Experimental Psychology: General*, 149(8): 1608–1613.
- Barrera, O.; Guriev, S.; Henry, E.; and Zhuravskaya, E. 2020. Facts, alternative facts, and fact checking in times of post-truth politics. *Journal of Public Economics*, 182: 104123.
- Bates, D.; Maechler, M.; Bolker, B.; Walker, S.; Christensen, R. H. B.; Singmann, H.; Dai, B.; Scheipl, F.; and Grothendieck, G. 2011. Package 'lme4'. *Linear mixed-effects models using S4 classes. R package version*, 1(6).
- Begg, I. M.; Anas, A.; and Farinacci, S. 1992. Dissociation of processes in belief: Source recollection, statement familiarity, and the illusion of truth. *Journal of Experimental Psychology: General*, 121(4): 446–458.
- Bode, L.; and Vraga, E. K. 2018. See something, say something: Correction of global health misinformation on social media. *Health Communication*, 33(9): 1131–1140.
- Bonilla, Y.; and Rosa, J. 2015. # Ferguson: Digital protest, hashtag ethnography, and the racial politics of social media in the United States. *American Ethnologist*, 42(1): 4–17.
- Brauer, M.; and Curtin, J. J. 2018. Linear mixed-effects models and the analysis of nonindependent data: A unified framework to analyze categorical and continuous independent variables that vary within-subjects and/or within-items. *Psychological Methods*, 23(3): 389–411.
- Braun, V.; and Clarke, V. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2): 77–101.
- Brown, S. C.; and Craik, F. I. 2000. Encoding and retrieval of information. *The Oxford handbook of memory*, 93–107.
- Carpenter, A. C.; and Schacter, D. L. 2017. Flexible retrieval: When true inferences produce false memories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(3): 335–349.
- Craik, F. I.; Govoni, R.; Naveh-Benjamin, M.; and Anderson, N. D. 1996. The effects of divided attention on encoding and retrieval processes in human memory. *Journal of Experimental Psychology: General*, 125(2): 159–180.
- Faul, F.; Erdfelder, E.; Lang, A.-G.; and Buchner, A. 2007. G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2): 175–191.

- Foster, J. L.; Huthwaite, T.; Yesberg, J. A.; Garry, M.; and Loftus, E. F. 2012. Repetition, not number of sources, increases both susceptibility to misinformation and confidence in the accuracy of eyewitnesses. *Acta Psychologica*, 139(2): 320–326.
- Frederick, S. 2005. Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4): 25–42.
- Funder, D. C.; and Ozer, D. J. 2019. Evaluating effect size in psychological research: Sense and nonsense. *Advances in Methods and Practices in Psychological Science*, 2(2): 156–168.
- Gelman, A. 2018. Discussion of effects of growth mindset: Let's not demand unrealistic effect sizes. *Statistical Modeling, Causal Inference, and Social Science*.
- Guess, A. M.; and Munger, K. 2022. Digital literacy and online political behavior. *Political Science Research and Methods*, 1–19.
- Hauser, D. J.; and Schwarz, N. 2016. Attentive Turkers: MTurk participants perform better on online attention checks than do subject pool participants. *Behavior Research Methods*, 48(1): 400–407.
- Huang, J.; Thornton, K. M.; and Efthimiadis, E. N. 2010. Conversational tagging in twitter. In *Proceedings of the 21st ACM conference on Hypertext and hypermedia*, 173–178.
- Hunt, R. R.; and Lamb, C. A. 2001. What causes the isolation effect? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27(6): 1359–1366.
- Keppel, G. 1991. *Design and analysis: A researcher's handbook*. NJ: Prentice-Hall, Inc.
- Lakens, D. 2013. Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. *Frontiers in Psychology*, 4.
- Lazer, D. M.; Baum, M. A.; Benkler, Y.; Berinsky, A. J.; Greenhill, K. M.; Menczer, F.; Metzger, M. J.; Nyhan, B.; Pennycook, G.; Rothschild, D.; et al. 2018. The science of fake news. *Science*, 359(6380): 1094–1096.
- Lee, S.; Forrest, J. P.; Strait, J.; Seo, H.; Lee, D.; and Xiong, A. 2020. Beyond Cognitive Ability: Susceptibility to Fake News Is Also Explained by Associative Inference. In *Extended Abstracts of the 2020 ACM CHI Conference*, 320:1–8.
- Lewandowsky, S.; and Van Der Linden, S. 2021. Countering misinformation and fake news through inoculation and prebunking. *European Review of Social Psychology*, 32(2): 348–384.
- Loftus, E. F. 2005. Planting misinformation in the human mind: A 30-year investigation of the malleability of memory. *Learning & Memory*, 12(4): 361–366.
- Loomba, S.; de Figueiredo, A.; Piatek, S. J.; de Graaf, K.; and Larson, H. J. 2021. Measuring the impact of COVID-19 vaccine misinformation on vaccination intent in the UK and USA. *Nature Human Behaviour*, 5(3): 337–348.
- Lorch, R. F. 1989. Text-signaling devices and their effects on reading and memory processes. *Educational psychology review*, 1(3): 209–234.
- Lutzke, L.; Drummond, C.; Slovic, P.; and Árvai, J. 2019. Priming critical thinking: Simple interventions limit the influence of fake news about climate change on Facebook. *Global Environmental Change*, 58: 101964.
- Malhotra, N.; Krosnick, J. A.; and Haertel, E. 2007. The psychometric properties of the GSS Wordsum vocabulary test. *GSS Methodological Report*, 11: 1–63.
- Murphy, G.; Loftus, E. F.; Grady, R. H.; Levine, L. J.; and Greene, C. M. 2019. False Memories for Fake News During Ireland's Abortion Referendum. *Psychological Science*, 30(10): 1449–1459.
- Pennycook, G.; Cannon, T. D.; and Rand, D. G. 2018. Prior exposure increases perceived accuracy of fake news. *Journal of Experimental Psychology: General*, 147(12): 1865–1880.
- Pennycook, G.; Epstein, Z.; Mosleh, M.; Arechar, A. A.; Eckles, D.; and Rand, D. G. 2021. Shifting attention to accuracy can reduce misinformation online. *Nature*, 592(7855): 590–595.
- Pennycook, G.; and Rand, D. G. 2019. Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition*, 188: 39–50.
- Perlow, J. 2020. Coronavirus misinformation spreading fast: Fake news on COVID-19 shared far more than CDC, WHO reports. <https://www.zdnet.com/article/coronavirus-misinformation-is-increasing-newsguard-finds/>. Accessed: 2022-05-14.
- Peterson, S. E. 1992. The cognitive functions of underlining as a study technique. *Literacy Research and Instruction*, 31(2): 49–56.
- Rho, E. H. R.; and Mazmanian, M. 2019. Hashtag Burnout? A Control Experiment Investigating How Political Hashtags Shape Reactions to News Content. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW): 1–25.
- Rights, J. D.; and Sterba, S. K. 2018. A framework of R-squared measures for single-level and multilevel regression mixture models. *Psychological Methods*, 23(3): 434–457.
- Rosenberg, S. 2016. EU Referendum: What Does Russia Gain from Brexit? *BBC News*.
- Seo, H.; Xiong, A.; and Lee, D. 2019. Trust It or Not: Effects of Machine-Learning Warnings in Helping Individuals Mitigate Misinformation. In *Proceedings of the 10th ACM Conference on Web Science*, 265–274.
- Shen, T. J.; Cowell, R.; Gupta, A.; Le, T.; Yadav, A.; and Lee, D. 2019. How gullible are you? Predicting susceptibility to fake news. In *Proceedings of the 10th ACM Conference on Web Science*, 287–288.
- Teng, X.; Lin, Y.-R.; Chung, W.-T.; Li, A.; and Kovashka, A. 2022. Characterizing User Susceptibility to COVID-19 Misinformation on Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, 1005–1016.
- Thomson, K. S.; and Oppenheimer, D. M. 2016. Investigating an alternate form of the cognitive reflection test. *Judgment and Decision Making*, 11(1): 99–113.
- Visentin, M.; Pizzi, G.; and Pichierri, M. 2019. Fake news, real problems for brands: The impact of content truthfulness and source credibility on consumers' behavioral intentions toward the advertised brands. *Journal of Interactive Marketing*, 45: 99–112.
- Vosoughi, S.; Roy, D.; and Aral, S. 2018. The spread of true and false news online. *Science*, 359(6380): 1146–1151.
- Wechsler, D. 2008. Wechsler adult intelligence scale—Fourth Edition (WAIS-IV). *San Antonio, TX: NCS Pearson*, 22(498): 816–827.
- Worthen, R. R. H. J. B.; et al. 2006. *Distinctiveness and memory*. Oxford University Press.
- Wu, L.; Morstatter, F.; Carley, K. M.; and Liu, H. 2019. Misinformation in social media: definition, manipulation, and detection. *ACM SIGKDD Explorations Newsletter*, 21(2): 80–90.
- Xiong, A.; Lee, S.; Seo, H.; and Lee, D. 2023. Effects of associative inference on individuals' susceptibility to misinformation. *Journal of Experimental Psychology: Applied*, 29(1): 1–17.
- Zeithamova, D.; and Preston, A. R. 2010. Flexible memories: differential roles for medial temporal lobe and prefrontal cortex in cross-episode binding. *Journal of Neuroscience*, 30(44): 14676–14684.