Reminders and Repetition of Misinformation: Helping or Hindering its Retraction?

Ullrich K. H. Ecker[1], Joshua L. Hogan[1], & Stephan Lewandowsky[2,1]

[1]School of Psychology, University of Western Australia, Australia

[2]School of Experimental Psychology, University of Bristol, UK

Word count: 4,621 (excluding title page, abstract, general summary, references, and online supplement)

Word count of Introduction and Discussion: 2,027

Corresponding Author: Ullrich Ecker, School of Psychology, University of Western Australia (M304), 35 Stirling Hwy, Perth 6009, Australia; telephone: +618 6488 3257;

e-mail: ullrich.ecker@uwa.edu.au

Abstract

People frequently rely on information even after it has been retracted, a phenomenon known as the continued-influence effect of misinformation. One factor proposed to explain the ineffectiveness of retractions is that repeating misinformation during a correction may inadvertently strengthen the misinformation by making it more familiar. Practitioners are therefore often encouraged to design corrections that avoid misinformation repetition. The current study tested this recommendation, investigating whether retractions become more or less effective when they include reminders or repetitions of the initial misinformation. Participants read fictional reports, some of which contained retractions of previous information, and inferential reasoning was measured via questionnaire. Retractions varied in the extent to which they served as misinformation reminders. Retractions that explicitly repeated the misinformation were more effective in reducing misinformation effects than retractions that avoided repetition, presumably because of enhanced salience. Recommendations for effective myth debunking may thus need to be revised.

Keywords: continued-influence effect; misinformation; myth debunking; familiarity

General Audience Summary

Information that is thought to be true but then turns out to be incorrect—so-called misinformation—can affect people's thinking and decision making even after it has been clearly corrected by a credible source, and even if people understand and later remember the correction. It has been proposed that one reason why corrections are so ineffective is that a myth is often repeated when it is corrected—explaining that vaccines do not cause autism almost necessarily repeats the association between vaccines and autism. This repetition can make the myth more familiar such that it comes to mind more easily in the future. Based on this notion, one recommendation to "myth debunkers" has been to avoid myth repetition in a correction. The present study directly tested this recommendation. We presented participants with news reports that did or did not contain corrections; these corrections did or did not repeat the to-be-corrected misinformation explicitly. We found—contrary to the popular recommendation—that corrections were more effective when they explicitly repeated the myth. Thus, it seems "safe" and even beneficial to repeat the myth explicitly when debunking it.

Reminders and Repetition of Misinformation: Helping or Hindering its Retraction?

Information that is initially presented as true but later identified as false and explicitly retracted often continues to influence people's cognition. This phenomenon is known as the continued-influence effect (CIE) of misinformation (H. Johnson & Seifert, 1994; Wilkes & Leatherbarrow, 1988). Research on the CIE has traditionally used a paradigm in which individuals read a (fictional) news report or scenario that includes a piece of critical information that subsequently is or is not retracted. The typical finding is that people's inferential reasoning, as for example measured through questionnaire, continues to be affected by the critical information despite clear and credible retractions, and even when individuals demonstrably understand and later remember the retraction (H. Johnson & Seifert, 1994, and Wilkes & Leatherbarrow, 1988; for reviews, see Lewandowsky, Ecker, Schwarz, Seifert, & Cook, 2012, and Seifert, 2002; for more recent work, see Ecker, Lewandowsky, E. P. Chang, & Pillai, R., 2014; Ecker, Lewandowsky, Cheung, & Maybery, 2015; Ecker, Lewandowsky, Fenton, & Martin, 2013; Guillory & Geraci, 2013, 2016; Nyhan, Reifler, & Ubel, 2013; Rich & Zaragoza, 2016; Thorson, 2016). In most of these studies, the retraction does have an effect—reliance on the critical information is typically halved compared to the no-retraction control—but the critical information almost always continues to be used to a significant extent.

Such continued reliance on misinformation is of particular concern when important decisions are at stake. One of the most commonly used examples of the CIE's real-world relevance is the ongoing impact of the fabricated link between childhood vaccines and autism, which has proven fairly resistant to correction (e.g., Poland & Spier, 2010). These real-world implications of the CIE are one of the factors that that have stimulated research effort into designing more effective correction strategies (cf. Cook & Lewandowsky, 2011; Lewandowsky et al., 2012; Schwarz, Newman, & Leach, in press).

One of the recommendations that has arisen from these efforts is to avoid repeating the misinformation when correcting it. This recommendation is founded in psychological theorizing that repeating the misinformation when retracting it may inadvertently strengthen the misinformation by making it more familiar. As it is well-known that familiar claims are more likely to be trusted and believed (e.g., Dechene, Stahl, Hansen, & Wanke, 2010; Weaver, Garcia, Schwarz, & Miller, 2007), the retraction could ironically backfire and increase reliance on misinformation rather than reduce it. Repeating the misinformation while identifying it as false could thus later leave people thinking "I've heard that before, so there's probably something to it" (Lewandowsky et al., 2012, p.115).

Some evidence for this "familiarity backfire effect" comes from a study by Skurnik, Yoon, and Schwarz (2007; also see Skurnik, Yoon, Park, & Schwarz, 2005), who provided participants with a "myths vs. facts" flyer that listed a number of claims regarding the flu vaccine, which were either affirmed or retracted. Skurnik et al. (2007) found that after a delay of 30 minutes, a substantial proportion of retracted myths were misremembered as facts, presumably based on the retraction-induced boost to the familiarity of the myths.[1]

More recently, Swire, Ecker, and Lewandowsky (2016) also investigated the role of familiarity in myth corrections. Participants were given a set of true and false claims of unclear

[1] In this study, the facts and myths all concerned the same topic, so an alternative account may involve source confusion (cf. M. Johnson, Hashtroudi, & Lindsay, 1993): participants may have just been confused about which statements were affirmed and which retracted. However, the effect was asymmetrical, in that a delay only led to increased acceptance of myths as true, with the rate of fact rejection remaining stable over time. This pattern is more in line with a familiarity-based explanation.

veracity (e.g., the fact that dogs can smell certain types of cancer, or the myth that playing Mozart can improve a baby's intelligence), which were subsequently repeated and then either affirmed or retracted. Claim belief was then measured after various retention intervals of up to three weeks. Swire et al. found that over time, the impact of myth retractions was less sustained than the impact of fact affirmations. This asymmetry was explained within a dual-processing framework, assuming that belief ratings can be based both on recollection of the affirmative/corrective explanation or the claim's familiarity (cf. Jacoby, 1991). The authors argued that for facts, it does not matter if belief is based on recollection of the affirmation or the familiarity of the claim—both will lead to acceptance of the fact; for myths, however, recollection of the retraction will lead to accurate rejection, whereas familiarity of the claim may lead to erroneous acceptance of the myth as true. The CIE thus seems at least partially familiarity-based. However, Swire et al. observed no familiarity backfire effect: myth belief post-retraction did not return to or exceed a pre-manipulation baseline (also see Peter & Koch, 2016). In sum, there is evidence for a role of familiarity in the CIE, but the evidentiary foundation for the recommendation that misinformation should not be repeated during its retraction is relatively weak.

Some theoretical accounts that focus on the *salience* of the misinformation during the correction even suggest that repeating misinformation when retracting it may be beneficial. Putnam, Wahlheim, and Jacoby (2014) as well as Stadtler, Scharrer, Brummernhenrich, and Bromme (2013) argued that detection of a conflict between rival event interpretations facilitates updating of a person's mental model of an event (cf. Morrow, Bower, & Greenspan, 1989). Such conflict detection is arguably more likely to occur if the retraction explicitly refers to both the invalidated interpretation as well as the new correct interpretation. Likewise, Kendeou, Walsh, Smith, and O'Brien (2014) argued that effective knowledge revision requires the co-activation of invalidated and correct event interpretations, which again is more likely to occur if the misinformation is explicitly repeated when it is retracted.

The current study

The current study aimed to determine whether providing reminders or repetitions of misinformation in the course of a retraction increased or decreased the subsequent CIE, thus testing the contrasting predictions of familiarity and salience accounts. In order to test these predictions, we presented participants with fictional news articles, some of which contained a retraction of earlier information, together with an alternative account of the respective event. The retraction either (1) did not refer back to the to-be-retracted misinformation, (2) included a reminder, explaining that the initial information was incorrect (without repeating the misinformation), or (3) explicitly repeated the misinformation before correcting it.

Method

The current study employed a within-subject design, featuring a single, four-level factor. The independent variable was the type of retraction condition. The dependent measure was participants' reliance on retracted misinformation, calculated based on responses to a questionnaire assessing participants' inferential reasoning.

Participants

A-priori power analysis suggested that to detect a small-to-medium difference between two conditions of effect size $f = 0.2$, with $\alpha = .05$ and $1 - \beta = .80$, and a moderate correlation between repeated measures of $r = .50$, the required sample size was 52 (this corresponds with the effect size found between conditions presenting misinformation once vs. thrice in Ecker, Lewandowsky, Swire, & D. Chang, 2011; power analysis was conducted with G*Power 3; Faul, Erdfelder, Lang, & Buchner, 2007). A total of $N = 60$ first year undergraduates from the University of Western Australia were recruited for participation in the current study, in return for partial course credit. The sample consisted of 18 male and 42 female participants, ranging from 17 to 53 years of age ($M = 20.52$, $SD = 7.14$).

Stimuli

Scenarios. Participants read six scenarios; they were informed that the scenarios would be the subject of a later memory test. Each scenario comprised two short articles and contained information regarding an unfolding news event (e.g., a wild fire). The first article in each case introduced the scenario and explained what happened; embedded in the first article was a piece of critical information that served as the potential target of a retraction in the second article (e.g., 'the fire had been deliberately lit'). The second article contained additional information pertaining to each of the scenarios; there were four versions of each second article, based on the type of retraction condition (see Online Supplement for all articles).

In the no-retraction (NR) control condition, the second article did not contain any retraction of information given in the first article. The other three conditions were retraction conditions. In the retraction-with-no-reminder (RNR) condition, more recent information given in the second article naturally superseded the initial misinformation account of the first article without any explicit reference to it (e.g., 'After a full investigation and review of witness reports, authorities have concluded that the fire was set off by lightning strikes'). The retraction-with-subtle reminder (RSR) condition contained a retraction featuring a subtle reminder of the initial account, explaining that it was incorrect (e.g., 'After a full investigation and review of witness reports, authorities have concluded that original reports were incorrect, and that the fire was set off by lightning strikes'). The final condition featured a correction that explicitly repeated the initial misinformation before retracting it (retraction-with-explicit-reminder condition, RER; e.g., ‘It was originally reported that the fire had been deliberately lit, but authorities have now ruled out this possibility. After a full investigation and review of witness reports, it has been concluded that the fire was set off by lightning strikes’).

Participants received three scenarios in the NR condition and one scenario in each of the three retraction conditions. We counterbalanced assignment of scenarios to conditions across participants, controlling presentation order such that (1) a no-retraction scenario was always presented first, (2) there were never two retractions presented consecutively, (3) each of the three retraction conditions occurred equally often at each of the three possible order positions 2, 4, and 6, and (4) each scenario occurred equally often at each order position. To this end, participants were randomly allocated to 1 of 6 pre-defined presentation orders of a Latin square design. This design was implemented in part to avoid participants being led to expect a retraction.

Participants read the six scenarios one-after-the-other, in the specified presentation order. The scenarios were presented via a slide-show on a computer screen. Participants read the first article and second article of each scenario on separate slides, before moving on to the next scenario. Each article was presented for a fixed amount of time (0.35 seconds per word), in order to control encoding time. This fixed time was pre-determined to allow reading times that were comfortable but not excessive. Participants were provided with a visual aid (a colored bar) on the screen that began to disappear slowly when there were 10 seconds left on the slide.

Questionnaire. We assessed participants’ understanding of the scenarios with a questionnaire (see Online Supplement). The questionnaire was presented in a booklet, following the order of scenarios established during study (specified by the pre-defined presentation order). The questionnaire comprised memory questions and inferential reasoning questions. For each scenario, participants' memory was assessed with an open-ended free recall question (e.g., 'Briefly summarize the 'wild-fire' article') and three multiple-choice questions with four possible alternatives (e.g., 'Where did the wild fire occur?'). These questions assessed adequate encoding and retention of scenario details.

Inferential reasoning questions required participants to make inferential judgments pertaining to the events in the scenarios. For each scenario, there were four open-ended questions designed to elicit responses relating to the critical information, while also allowing participants the opportunity to cite unrelated, alternative responses (e.g., 'How could such events be prevented in future?'). In addition, there were three rating-scale questions requiring participants to indicate on a 10-point scale their level of agreement with a statement (e.g., 'Would it be lawful for someone to be punished as a result of the wild fire?').

Procedure

Participants read an ethically-approved information sheet and provided informed consent. Participants then read the six scenarios in individual testing booths. After readings the scenarios, participants completed an unrelated distractor task for approximately 30 minutes, following Skurnik et al. (2007). Finally, participants completed the questionnaire assessing their understanding of the scenarios. The entire experiment took approximately one hour to complete.

Results

Questionnaire scoring

Questionnaire responses were coded by a scorer who was blind to experimental condition, following a standardized guide.

Memory scores. Recall of several aspects of the scenarios was scored separately; in particular, there were scores for (a) general fact-recall of arbitrary details, (b) recall of the critical information, (c) recall of the retraction, and (d) recall of the alternative.

The general fact-recall score was calculated based on responses to both the open-ended free recall question and the multiple-choice questions. Scoring of the free recall item was based on predetermined ‘idea units’. Idea units pertained to information contained in the scenarios that did not refer to the critical information or its alternative, and that was not assessed by the multiple-choice questions. For each scenario, two major idea units (i.e., information considered a major theme of the scenario; e.g., that the wild fire had not caused damage to residential property) and two minor idea units (i.e., information considered a minor detail in the scenario; e.g., that the wild fire had damaged forest reserves) were identified a priori (see Online Supplement for all idea units). A score of 1 was given for recall of a major idea unit, while a score of 0.5 was given for recall of a minor idea unit, resulting in a possible maximum recall score of 3 for each scenario. Additionally, correct responses to multiple-choice questions were given a score of 1, resulting in a possible maximum score of 3 for each scenario. Scores were then combined and scaled to yield a final memory score for each scenario ranging from 0 to 1. The memory scores of the three non-retraction scenarios were collapsed, such that each participant had one memory score per experimental condition.

Memory for the critical piece of information, memory for the retraction, and memory for the alternative account was coded in separate scores based on the response to the open-ended free-recall question. For each scenario, the score was 1 when the respective piece of information (i.e., the critical information, the retraction, or the alternative) was recalled and 0 otherwise. To illustrate, this means that any mention of the critical information led to a critical-information recall score of 1, whether or not the participant concurrently or subsequently mentioned the retraction (e.g., in the fire scenario, “it was thought the fire was caused by arson” and “the fire was not caused by arson as initially thought” were both scored 1 for critical-information recall, with the latter also receiving a retraction-recall score of 1). This means that recalling the initial critical piece of information does not necessarily imply reliance on misinformation, as long as a participant also recalled the retraction and/or alternative. Also, it was possible that the retraction would be recalled without mention of the critical information (e.g., “initial speculations were not confirmed”). Finally, any mention of the alternative led to an alternative-recall score of 1, irrespective of whether a retraction was mentioned (e.g., “lightning caused the fire” or “initial speculations were not confirmed, and it was concluded the fire was caused by lightning” both led to an alternative-recall score of 1, with the latter also scoring a 1 for retraction recall). It was possible that all three measures were scored 1 (e.g., “the fire was not caused by arson as initially thought but by lightning”). Retraction and alternative recall scores were not coded for the NR condition.

Inferential reasoning scores. For each scenario, an inference score was calculated based on responses to the four open-ended inference questions and the three rating scales. For each open-ended question, a score of 1 was awarded for a clear and uncontroverted reference to the critical information (e.g., an answer such as “Arson” in response to the question “What was the cause of the fire?”). A score of 0 was given for any other response (e.g., a controverted answer such as “It was initially thought it was arson, but that was not true”). Rating-scale scores ranged from 1 to 10, with higher scores denoting stronger reliance on the critical information (scales that were negatively worded to this end were reverse-scored). For each scenario, all seven question scores were equally weighted, combined, and transformed into an inference score ranging from 0 to 1. The inference scores of the three non-retraction scenarios were collapsed, such that each participant had one inference score per experimental condition.

Analysis

Preliminary analyses were conducted to determine whether any participants needed to be removed from further analysis. The fact-recall scores were examined to determine whether any participants scored lower than an a-priori criterion of 0.167 (1 out of the maximum of 6) for all scenarios. One participant violated this, but as they scored above the criterion in 5 of the 6 scenarios, their data were retained, and thus no participants were excluded based on this criterion.[2] The data were then screened for outliers, but none were identified.

Memory scores. Memory scores are provided in the top row of Table 1. Scores were analyzed to investigate whether there were any differences between conditions in comprehension of and memory for the scenarios. The mean memory scores across conditions were comparable, and a one-way repeated-measures analysis of variance (ANOVA) revealed no significant effect of condition, although the analysis just missed the conventional significance criterion $F(3,177) = 2.35$, $p = .07$, $\eta_p^2 = .04$.

Scores on critical-information recall and alternative recall were also comparable across conditions (see Table 1). Non-parametric repeated measures ANOVAs (Friedman tests) found no significant differences, $\chi^2 < 1$ for critical-information recall, and $\chi^2(2) = 3.60$, $p = .17$, for alternative recall. Next, retraction recall was analyzed to determine whether there were any differences in recall of the retraction between conditions. Mean retraction recall scores are also given in Table 1. A Friedman test revealed a significant main effect of condition on retraction recall, $\chi^2(2) = 7.28$, $p = .03$. A contrast analysis revealed a significant difference between the

[2] All analyses were repeated without this participant; this did not affect results.

RNR and RER conditions, $\chi^2(1) = 8.07$, $p < .01.04$, but not between RNR and RSR, $\chi^2(1) = 1.67$, $p = .20$, or RSR and RER, $\chi^2(1) = 1.80$, $p = .18$.

As an initial test of the question if reliance on misinformation differed between retraction conditions, we calculated a measure of misinformation reliance by simply subtracting the summed retraction-recall and alternative-recall scores from the critical-information recall score, separately for each retraction condition. This misinformation score was 1 if and only if the critical misinformation was recalled without the retraction or the alternative being recalled as well; if the misinformation was not recalled, or if it was recalled alongside its retraction and/or the alternative, the score was 0 or -1 (a score of -2 was theoretically possible but did not eventuate). Thus, more reliance on misinformation is reflected in more positive scores. The mean misinformation scores across conditions are given in the bottom row of Table 1. A Friedman test yielded a significant main effect of condition, $\chi^2(2) = 6.57$, $p = .04$, substantiating that misinformation reliance was greatest in the RNR and lowest in the RER condition [in a contrast analysis, the RNR-RER difference was significant, $\chi^2(1) = 5.77$, $p = .02$, but the RNR-RSR and RSR-RER differences were not, $\chi^2(1) < 3.21$, $p > .12$].

Table 1

Memory, Recall, and Misinformation Scores Across Conditions

	NR		RNR		RSR		RER	
	M	*SE*	*M*	*SE*	*M*	*SE*	*M*	*SE*
Memory score	0.66	0.01	0.62	0.02	0.65	0.02	0.62	0.02
Critical-information recall	0.53	0.04	0.53	0.06	0.50	0.07	0.53	0.06
Alternative recall	-	-	0.33	0.06	0.43	0.06	0.48	0.07
Retraction recall	-	-	0.13	0.04	0.22	0.05	0.32	0.06
Misinformation score			0.07	0.10	-0.15	0.10	-0.27	0.10

Note. NR, no-retraction condition; RNR, retraction with no reminder condition; RSR, retraction with subtle reminder condition; RER, retraction with explicit reminder condition.

Inferential reasoning scores. The mean inference scores are depicted in Figure 1; mean scores were M_{NR} = 0.58 (SE_{NR} = 0.02), M_{RNR} = 0.39 (SE_{RNR} = 0.03), M_{RSR} = 0.34 (SE_{RSR} = 0.03), and M_{RER} = 0.27 (SE_{RER} = 0.03). First, one-sample *t*-tests were conducted to determine whether inference scores differed significantly from zero (zero representing no reliance on misinformation in reasoning). Results revealed that inference scores were substantially greater than zero in all retraction conditions, all $ts(59) > 9.96$, $p < .001$, indicating presence of a CIE in all three retraction conditions.

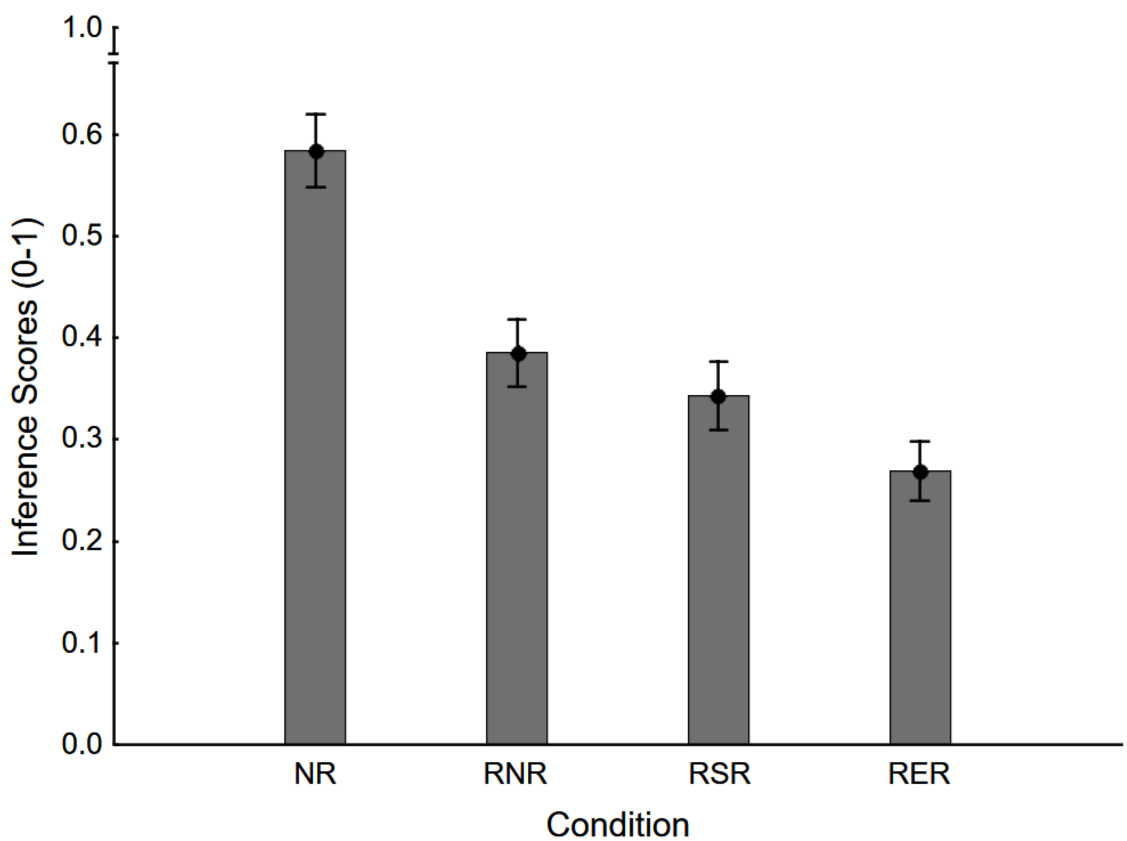

Figure 1. Mean inference scores (0-1) across experimental conditions; NR = no retraction; RNR = retraction with no reminder; RSR = retraction with subtle reminder; RER = retraction with explicit reminder. Error bars depict within-subject standard errors of the mean. See text for details.

A repeated-measures ANOVA on inference scores revealed a significant main effect of retraction condition, $F(3,177) = 22.24$, $p < .001$, $\eta_p^2 = .27$. A contrast analysis (see Table 2) revealed pairwise significant differences between the NR and all three retraction conditions, indicating reduced reliance on the critical information after any type of retraction. There was also a significant difference between RNR and RER conditions, indicating that reliance on misinformation was further reduced if the retraction featured an explicit repetition of the misinformation. The difference between the RSR and RER conditions was marginally significant, suggesting that an explicit reminder tended to make the retraction more effective than a subtle reminder.[3]

Table 2

Contrasts on Inference Scores

Contrast	$F(1,59)$	p
NR vs. RNR	18.83	< .001
NR vs. RSR	30.21	< .001
NR vs. RER	73.99	< .001
RNR vs. RSR	1.26	.27
RNR vs. RER	9.60	.003
RSR vs. RER	3.44	.07

Note. NR = no retraction; RNR = retraction with no reminder; RSR = retraction with subtle reminder; RER = retraction with explicit reminder.

Discussion

This study investigated whether providing reminders or repetitions of misinformation in the course of a retraction increased or decreased people's reliance on misinformation. In doing so, the current study compared contrasting theoretical predictions of the familiarity account of

[3] Observed effects naturally differed somewhat across the different scenarios. Figure 1 of the online supplement shows the data across scenarios. However, given the low number of observations per cell after splitting up the scenarios, we caution against over-interpretation of these differences.

continued-influence effects and salience accounts of mental-model updating and knowledge revision, and directly addressed the common recommendation to not repeat the misinformation during its retraction in order to avoid a familiarity backfire effect. To the best of our knowledge, the only previous test of the effect of an explicit repetition of the misinformation in the retraction was performed by Wilkes and Leatherbarrow (1988), in the study that first demonstrated the CIE. Wilkes and Leatherbarrow reported a null effect of the repetition (although the retraction was numerically more effective if the misinformation was repeated), but given the limited power of that study, no firm conclusions could be drawn.

We found that any kind of retraction reduced reliance on misinformation compared to a control condition with no retraction. In the current study, each retraction was accompanied by an alternative causal account of the event, which is a factor known to reliably reduce reliance on misinformation (e.g., Ecker, Lewandowsky, & Tang, 2010; H. Johnson & Seifert, 1994; Seifert, 2002). Nonetheless, misinformation continued to have a significant influence on inferential reasoning despite a retraction. In line with previous research, this indicates that the provision of an alternative account alone is not enough to eliminate the CIE (e.g., Ecker, Lewandowsky, & Apai, 2011; Ecker, Lewandowsky, Swire, & D. Chang, 2011; Guillory & Geraci, 2013; 2016; Wilkes & Leatherbarrow, 1988).

Notably, we found that a retraction featuring an explicit reminder was most effective in reducing reliance on misinformation. That is to say, a retraction that explicitly repeated the misinformation (condition RER) lowered reliance on misinformation more than a retraction that provided no reminder (condition RNR); it also tended to be more effective than a retraction providing only a subtle reminder merely pointing out that an earlier account was incorrect (condition RSR). In line with salience accounts of mental-model updating and knowledge revision (Kendeou et al., 2014; Putnam et al., 2014; Stadtler et al., 2013), we propose that an explicit-reminder retraction was more effective because it made both the falsity of the misinformation and the conflict between the outdated and updated event representations salient. Such a salient retraction may lead to stronger updating of a person's mental model of the event immediately upon encoding the retraction, which implies that the corrective effect of the retraction happens mainly at encoding rather than during the later retrieval phase at test. An alternative view that is congruent with our data is that the salience of the RER retractions made them more memorable and thus more potent at the time of retrieval (in line with retrieval accounts of the CIE, cf. Ecker et al., 2010; Seifert, 2002; Swire et al., 2016).

The relevance of this finding for application is clear-cut. Despite the known relevance of familiarity for the CIE (Peter & Koch, 2016; Swire et al., 2016; also Skurnik et al., 2007, but see Cameron et al., 2013, for a conceptual-replication failure), the common recommendation to avoid repetition of misinformation in the course of its retraction (Cook & Lewandowsky, 2011; Lewandowsky et al., 2012) needs to be qualified. Based on the current data, it seems that it may in fact be beneficial to repeat the to-be-corrected misinformation once when retracting it, to enhance the salience of the correction. Of course, familiarity can in principle still lead to myths being remembered as facts, and thus unnecessary repetition of misinformation should still be avoided, in particular in situations where one may familiarize people with misinformation they have not encountered before (Schwarz et al., in press). This qualified recommendation meshes well with findings from the educational literature, where it has been found that misinformation can be used as a teaching tool (Bedford, 2010; Cook, Bedford, & Mandia, 2014; Kowalski & Taylor, 2009; Osborne, 2010).

Before concluding that repeating misinformation in order to then refute it can be considered an effective and 'safe' strategy, we must address a number of caveats and limitations of the present study. First, the interval between presentation of misinformation and its retraction was shorter than it would be in many real-world situations. While there is little evidence that the duration of this interval has a direct impact on a retraction's effectiveness (Ecker et al., 2015, Experiment 1; H. Johnson & Seifert, 1994), the intervals examined in prior research did not differ much, and it is unknown how a longer misinformation-retraction interval might moderate the effect of misinformation repetition. Second, retractions in the present study featured causal alternatives, which in the real world are often not available. It is possible that misinformation repetition might have a negative effect only in the absence of causal alternatives (however, Swire et al., 2016, failed to find familiarity backfire effects even in the absence of alternatives). Finally, while the present study used fictional events, real-world myths are often part of pre-existing 'knowledge', and myth repetition may have a negative effect only in the latter case. However, this seems unlikely as repetition of an unfamiliar myth (as in the present study) will have a larger impact on the myth's relative familiarity compared to the repetition of an already-familiar myth. Thus, unfamiliar myths should make it more, not less, likely to detect a negative influence of repetition. In line with this, Swire et al. (2016) failed to find any evidence for familiarity backfire effects resulting from repetition of moderately familiar real-world myths.

To conclude, we presented evidence that repeating misinformation in the course of its retraction can reduce continued-influence effects. However, the influence of misinformation persisted despite the availability of causal alternatives and the repetition-enhanced effectiveness of retractions.

References

Ayers, M. S., & Reder, L. M. (1998). A theoretical review of the misinformation effect: Predictions from an activation-based memory model. *Psychonomic Bulletin & Review*, *5*, 1-21.

Bedford, D. (2010). Agnotology as a teaching tool: Learning climate science by studying misinformation. *Journal of Geography, 109*, 159-165.

Cameron, K. A., Roloff, M. E., Friesema, E. M., Brown, T., Jovanovic, B. D., Hauber, S., & Baker, D. W. (2013). Patient knowledge and recall of health information following exposure to "facts and myths" message format variations. *Patient Education and Counseling, 92*, 381-387.

Cook, J., Bedford, D., & Mandia, S. (2014). Raising climate literacy through addressing misinformation: Case studies in agnotology-based learning. *Journal of Geoscience Education, 62*, 296-306.

Cook, J., & Lewandowsky, S. (2011). *The Debunking Handbook*. St. Lucia, Australia: University of Queensland.

Dechene, A., Stahl, C., Hansen, J., & Wanke, M. (2010). The truth about the truth: A meta-analytic review of the truth effect. *Personality and Social Psychology Review, 14*, 238-257.

Ecker, U. K. H., Lewandowsky, S., & Apai, J. (2011a). Terrorists brought down the plane! No, actually it was a technical fault: Processing corrections of emotive information. *Quarterly Journal of Experimental Psychology*, *64*, 283-310.

Ecker, U. K. H., Lewandowsky, S., Chang, E. P., & Pillai, R. (2014). The effects of subtle misinformation in news headlines. *Journal of Experimental Psychology: Applied*, *20*, 323-335.

Ecker, U. K. H., Lewandowsky, S., Cheung, C. S. C., & Maybery, M. T. (2015). He did it! She did it! No, she did not! Multiple causal explanations and the continued influence of misinformation. *Journal of Memory and Language*, *85*,101-115.

Ecker, U. K. H., Lewandowsky, S., Fenton, O., & Martin, K. (2014). Do people keep believing because they want to? Pre-existing attitudes and the continued influence of misinformation. *Memory & Cognition*, *42*, 292-304.

Ecker, U. K. H., Lewandowsky, S., Swire, B., & Chang, D. (2011b). Correcting false information in memory: Manipulating the strength of misinformation encoding and its retraction. *Psychonomic Bulletin & Review*, *18*, 570-578.

Ecker, U. K. H., Lewandowsky, S., & Tang, D. (2010). Explicit warnings reduce but do not eliminate the continued influence of misinformation. *Memory & Cognition*, *38*, 1087-1100.

Guillory, J. J., & Geraci, L. (2013). Correcting erroneous inferences in memory: The role of source credibility. *Journal of Applied Research in Memory and Cognition*, *2*, 201-209.

Guillory, J. J., & Geraci, L. (2016). The persistence of erroneous information in memory: The effect of valence on the acceptance of corrected information. *Applied Cognitive Psychology*, *30*, 282-288.

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*, 513-541.

Johnson, H. M., & Seifert, C. M. (1994). Sources of the continued influence effect: When misinformation in memory affects later inferences. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *20*, 1420-1436.

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin*, *114*, 3-28.

Kendeou, P., Walsh, E. K., Smith, E. R., & O'Brien, E. J. (2014). Knowledge revision processes in refutation texts. *Discourse Processes*, *51*, 374-397.

Kowalski, P., & Taylor, A. K. (2009). The effect of refuting misconceptions in the introductory psychology class. *Teaching of Psychology, 36*, 153-159.

Lewandowsky, S., Ecker, U. K. H., Schwarz, N., Seifert, C., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest*, *13*, 106-131.

Morey, R. D. (2008). Confidence intervals from normalized data: A correction to Cousineau (2005). *Tutorial in Quantitative Methods for Psychology, 4*, 61-64.

Morrow, D. G., Bower, G. H., & Greenspan, S. L. (1989). Updating situation models during narrative comprehension. *Journal of Memory and Language*, *28*, 292-312.

Nyhan, B., Reifler, J., & Ubel, P. A. (2013). The hazards of correcting myths about health care reform. *Medical Care*, *51*, 127-132.

Osborne, J. (2010). Arguing to learn in science: The role of collaborative, critical discourse. *Science, 328*, 463-466.

Peter, C., & Koch, T. (2016). When debunking scientific myths fails (and when it does not): The backfire effect in the context of journalistic coverage and immediate judgments as prevention strategy. *Science Communication, 38,* 3-25.

Poland, G. A., & Spier, R. (2010). Fear, misinformation, and innumerates: How the Wakefield paper, the press, and advocacy groups damaged the public health. *Vaccine, 28*, 2361-2362.

Putnam, A. L., Wahlheim, C. N., & Jacoby, L. L. (2014). Memory for flip-flopping: Detection and recollection of political contradictions. *Memory & Cognition, 42*, 1198-1210.

Rich, P. R., & Zaragoza, M. S. (2016). The continued influence of implied and explicitly stated misinformation in news reports. *Journal of Experimental Psychology: Learning Memory and Cognition, 42*, 62-74.

Schwarz, N., Newman, E., & Leach, W. (in press). Making the truth stick and the myths fade: Lessons from cognitive psychology. *Behavioural Science and Policy.*

Seifert, C. M. (2002). The continued influence of misinformation in memory: What makes a correction effective? *The Psychology of Learning and Motivation, 41*, 265-292.

Skurnik, I., Yoon, C., Park, D. C., & Schwarz, N. (2005). How warnings about false claims become recommendations. *Journal of Consumer Research, 31*, 713-724.

Skurnik, I., Yoon, C., & Schwarz, N. (2007). *Education about flu can reduce intentions to get a vaccination.* Unpublished manuscript.

Stadtler, M., Scharrer, L., Brummernhenrich, B., & Bromme, R. (2013). Dealing with uncertainty: Readers' memory for and use of conflicting information from science texts as function of presentation format and source expertise. *Cognition & Instruction, 31*, 130-150.

Swire, B., Ecker, U. K. H., & Lewandowsky, S. (2016). *The role of familiarity in correcting inaccurate information.* Manuscript submitted for publication.

Thorson, E. (2016). Belief echoes: The persistent effects of corrected misinformation. *Political Communication, 33*, 460-480.

Weaver, K., Garcia, S. M., Schwarz, N., & Miller, D. T. (2007). Inferring the popularity of an opinion from its familiarity: A repetitive voice can sound like a chorus. *Journal of Personality and Social Psychology, 92*, 821-833.

Wilkes, A. L., & Leatherbarrow, M. (1988). Editing episodic memory following the identification of error. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 40*, 361-387.

Author contributions: UE and SL conceptualized the study; UE and JH designed the study; JH and UE prepared the materials; JH collected the data; JH and UE analyzed the data; UE, JH, and SL wrote the manuscript.

Acknowledgements. This research was supported by a Discovery Grant from the Australian Research Council to the first and third author (DP160103596). We thank Charles Hanich for research assistance. The lab web address is http://www.cogsciwa.com.