

Correcting false information in memory: Manipulating the strength of misinformation encoding and its retraction

Ullrich K. H. Ecker · Stephan Lewandowsky ·
Briony Swire · Darren Chang

Published online: 26 February 2011
© Psychonomic Society, Inc. 2011

Abstract Information that is presumed to be true at encoding but later on turns out to be false (i.e., misinformation) often continues to influence memory and reasoning. In the present study, we investigated how the strength of encoding and the strength of a later retraction of the misinformation affect this continued influence effect. Participants read an event report containing misinformation and a subsequent correction. Encoding strength of the misinformation and correction were orthogonally manipulated either via repetition (Experiment 1) or by imposing a cognitive load during reading (Experiment 2). Results suggest that stronger retractions are effective in reducing the continued influence effects associated with strong misinformation encoding, but that even strong retractions fail to eliminate continued influence effects associated with relatively weak encoding. We present a simple computational model based on random sampling that captures this effect pattern, and conclude that the continued influence effect seems to defy most attempts to eliminate it.

Once encoded, information may continue to influence reasoning, even if it later turns out to be incorrect. The persistent reliance on such misinformation, even when people can recall a correction or retraction, has been labeled the *continued influence effect* (Johnson & Seifert, 1994). For example, if a fictional character is accused of a crime but is later exonerated, people continue to use the outdated misinformation (that the person is guilty) in subsequent inferences, even if they recall the correction. The continued influence effect

has been demonstrated in many settings (e.g., Ecker, Lewandowsky, & Apai, 2011; Ecker, Lewandowsky, & Tang, 2010; Johnson & Seifert, 1994; van Oostendorp, 1996; Wilkes & Leatherbarrow, 1988). In the real world, continued belief in unsubstantiated claims can have serious implications, as in the case of the purported link between certain vaccines and autism (Baron-Cohen, 2009) or between Iraq and weapons of mass destruction (WMDs; Kull, Ramsay, & Lewis, 2003; Lewandowsky, Stritzke, Oberauer, & Morales, 2005, 2009).

Two different approaches have been put forward to explain the continued influence effect. One of these refers to “mental event models” (Johnson & Seifert, 1994; Wilkes & Leatherbarrow, 1988). People are thought to build mental models of unfolding events, but seem reluctant to dismiss key information in their model (e.g., what caused an event) when no plausible alternative exists to fill the void. Accordingly, the provision of alternative causal information (e.g., presentation of an alternative suspect) has long been the only factor known to reduce the continued influence effect (Johnson & Seifert, 1994). When no alternative is presented, people prefer an inconsistent event model to an incomplete event model. Hence, in their inferential reasoning, they may rely on outdated information despite knowing that it is false, rather than acknowledging the lack of valid information available.

An alternative account of the continued influence effect can be formulated within dual-process theory (Ayers & Reder, 1998; Schwarz, Sanna, Skurnik, & Yoon, 2007). Ayers and Reder suggested that pieces of both valid and invalid information compete for automatic activation in memory. By contrast, recall of specifics (such as the source or the validity of the information) relies on strategic retrieval. The assumption thus is that memory for the retraction is based mainly on a controlled retrieval process

U. K. H. Ecker (✉) · S. Lewandowsky · B. Swire · D. Chang
School of Psychology, University of Western Australia,
Crawley, WA 6009, Australia
e-mail: ullrich.ecker@uwa.edu.au

that aims to integrate the available information in order to produce valid inferences. Continued influence arises when (a) the misinformation is supplied by an automatic retrieval process, whose output is mainly determined by memory strength, and (b) the strategic retrieval process fails, either because the person does not engage in strategic retrieval (e.g., inadequate output monitoring under time pressure; cf. Jacoby, 1999), or because of genuine failure of the process.

The understanding of misinformation effects outside the laboratory is complicated by the fact that both misinformation and its retraction are often disseminated repeatedly and/or with varying rigor. To use a notorious real-world example, the Bush administration purportedly made 935 false statements about the security risk posed by Iraq in the 2 years following 9/11 (Lewis & Reading-Smith, 2008). It is possible that the reiteration of this misinformation (i.e., that Iraq possessed WMDs) led to particularly powerful continued influence (e.g., the widespread continued belief in the existence of WMDs in Iraq; Kull et al., 2003; Lewandowsky et al., 2005, 2009). However, it is unclear what the effects of less extreme strength manipulations are when applied to the encoding of misinformation and/or the encoding of its retraction.

According to the event model approach, the initial integration of information into the event model is more readily performed than is its updating after a retraction. This is because updating involves additional processing: Not only does the retraction itself need to be encoded, but the existing misinformation (e.g., *X caused Y*) also needs to be removed from the event model, and the new information (e.g., *unclear what caused Y but it was “not X”*) must be integrated (cf. Ecker, Lewandowsky, Oberauer, & Chee, 2010; Radvansky & Copeland, 2001). This suggests that the retraction may profit more from repetition than does the encoding of misinformation; hence, we may expect differential effects of strengthening.

In contrast, the dual-processing account would predict the opposite. Eakin, Schreiber, and Sergent-Marshall (2003) demonstrated that misinformation effects could be suppressed by explicit warnings, which foster strategic monitoring processes (cf. Ecker et al., 2010). This suggests that the effects of retractions are mainly carried by strategic processing. However, warnings failed to reduce misinformation effects when misinformation was presented repeatedly, which, according to Eakin et al., mainly fostered automatic retrieval. On the assumption that strengthening information encoding will mainly affect automatic processes, repetition may have a greater impact on the encoding of misinformation than on its retraction.

The present study is the first to systematically investigate how the strength with which misinformation is encoded, and the vigor with which it is later retracted, affects the continued influence effect. It is well documented that

repetition enhances belief in the truth of repeated assertions (e.g., Allport & Lepkin, 1945; Weaver, Garcia, Schwarz, & Miller, 2007) as well as memory more generally. This is especially true if repetition occurs with some temporal spacing or in different contexts (cf. Chabot, Miller, & Juola, 1976; Verkoeijen, Rikers, & Schmidt, 2004). This in turn suggests that repetition unfolds its effects by associating information with various contexts or sources, which could serve as retrieval cues and/or lend more credibility to a repeatedly encoded piece of information. Hence, enhanced encoding could increase the continued influence effect because repeated misinformation may become harder to retract (cf. Schul & Mazursky, 1990).

Alternatively, however, it has been suggested that enhanced encoding of misinformation may *reduce* its continued influence, since memory updating may be more efficient when the initial information—despite being false—is well represented and active in memory (van Oostendorp, 1996). The idea that only something that is well represented in memory can be easily updated is in line with at least two related areas of enquiry. First, reconsolidation theory claims that information needs to be activated in order to be updated (e.g., Hupbach, Gomez, Hardt, & Nadel, 2007). Second, in the categorization and problem solving literature, it has become clear that knowing something well is no barrier for, or may even benefit, knowledge restructuring (i.e., shift to an alternate strategy; Sewell & Lewandowsky, 2011). Finally, in terms of dual processes, repetition will not only enhance automatic retrieval (as suggested by Eakin et al., 2003), but will usually lead to improved controlled memory processes as well, in particular improved source memory (e.g., Jacoby, 1999). Hence, inasmuch as factors such as source confusion are reduced by repetition, continued influence could also be reduced by repetition of misinformation.

Concerning the retraction, intuitively, if a statement is retracted with greater emphasis, one might expect less continued influence. However, the only study known to us that manipulated the strength of the retraction (van Oostendorp & Bonebakker, 1999) found continued influence to be unaffected by repetition of a retraction: Two retractions were found to be as (in-)effective as one. Moreover, the literature on metacognitive effects of repetition has shown that, ironically, because misinformation is often repeated when it is retracted, more frequent retraction of misinformation can paradoxically enhance its impact even after relatively short retention intervals of 30 min (Schwarz et al., 2007; Skurnik, Yoon, Park, & Schwarz, 2005). In other words, the retraction could serve as a recursive reminder of the misinformation (Hintzman, 2010). Such backfire effects of retractions have been observed in examinations of the effects of retractions on political misperceptions (Nyhan & Reifler, 2010) and mock

juror behavior (Pickel, 1995), and are obviously a reason for concern.

In summary, it is unclear whether continued-influence effects are necessarily enhanced and reduced, respectively, by repetition or by other means of strengthening the initial encoding of misinformation or its retraction.

We present two experiments that manipulated the strength with which misinformation was encoded and retracted. In Experiment 1, we orthogonally varied the number of repetitions of the misinformation and its retraction; in Experiment 2, we used a cognitive load manipulation at the encoding or the retraction stage. In both experiments, participants received an adaptation of a much-used warehouse fire script (Johnson & Seifert, 1994; Wilkes & Leatherbarrow, 1988), in which a fire was initially reported to have been caused by volatile materials stored negligently in a closet, and with subsequent reports retracting this, stating that the closet had been empty.

Experiment 1

Two between-subjects factors were fully crossed. Factors were the strength of misinformation (one or three repetitions), and the strength of retraction (zero, one, or three repetitions). Additionally, a control group with no mention of volatile materials was tested.

Method

Participants A total of 161 undergraduates from the University of Western Australia (108 females) participated and were randomly assigned to conditions¹ ($N = 23$ per condition).

Stimuli Participants received 17 messages, each printed on a separate page. The 0-MI control condition featured no statements referring to volatile materials, and obviously no retraction either. In the 1- and 3-MI conditions, a statement regarding the presence of volatile materials appeared once (Message 6) or three times (Messages 6, 7, and 8), respectively. In the 3-R conditions, the retraction was repeated three times (Messages 13, 14, and 15) as opposed to only once in the 1-R conditions (Message 14), and there was no retraction at all in the 0-R conditions. Repetitions were presented in different contexts (e.g., as a radio transmission from a police investigator/as information

passed on to the fire captain/in a public radio announcement). This was done to make the repetitions appear more natural and to enhance their potential impact by increasing contextual variation (e.g., Verkoeijen et al., 2004). Scripts were of equal length in all conditions; filler messages were added where needed.

Procedure Participants read the messages aloud at their own pace, without backtracking. After an unrelated 10-min distractor task, participants received an open-ended questionnaire, consisting of 10 causal inference questions (e.g., *What could have caused the explosions?*), 10 fact questions (e.g., *What time was the fire eventually put out?*), and two manipulation-check questions, targeting awareness of the retraction (e.g., *Was any of the information in the story subsequently corrected or altered? And if so, what was it?*), always administered in this order.

Results

Analysis focused on three dependent measures: the number of references to misinformation (i.e., the inference score), the accuracy of recall, and acknowledgment of the retraction. References to misinformation (viz. negligently stored volatile materials) were counted only if they were causal and uncontroverted.

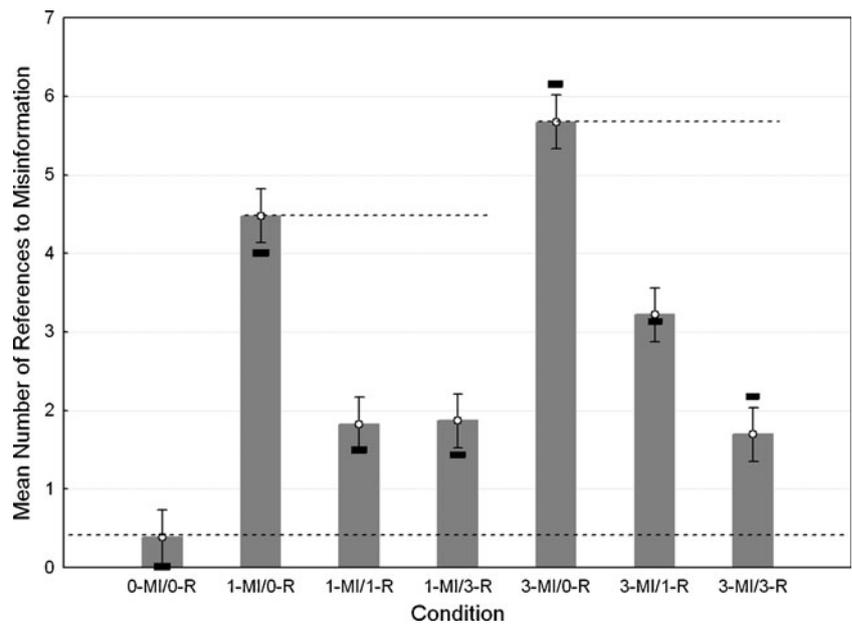
Coding procedure Responses were tallied by a naive scorer following a scoring guide. Inter-rater reliability with a second scorer was high ($r_s = .97, .82,$ and $.94$ for fact-recall, inference, and manipulation-check scores, respectively, based on a sample of 18 questionnaires).

Inferences Mean inference scores are shown in Fig. 1. The 0-MI and 0-R conditions provide empirical baselines for interpretation of the remaining experimental conditions and are represented by the dotted lines. The 0-MI control condition expectedly yielded few spontaneous references to misinformation (significantly fewer than those in the 1-MI/3-R condition, which was expected to have the lowest score of the remaining conditions; all contrasts are given in Table 1).

A two-way ANOVA on the six experimental conditions yielded significant main effects of strength of misinformation, $F(1, 132) = 7.32, p < .01, \eta^2 = .05,$ and strength of retraction, $F(2, 132) = 45.02, p < .001, \eta^2 = .41,$ which were qualified by a marginally significant interaction, $F(2, 132) = 2.74, p = .07, \eta^2 = .04.$ Planned contrasts (cf. Table 1) demonstrated that, not surprisingly, repeated misinformation encoding led to stronger misinformation effects when there was no or only one retraction (contrasts 4 and 8). After three presentations of misinformation, one retraction reduced reliance on misinformation (contrast 5),

¹ The 3-MI/0-R conditions of Experiment 1, as well as the L-MI/0-R conditions of Experiment 2 were tested after the other conditions, but testing was carried out in the same lab, by the same experimenter, and during the same time of year; participants were taken from the same pool.

Fig. 1 Mean number of references to misinformation for all conditions in Experiment 1; error bars show standard errors of the mean. Black rectangles indicate predictions of the sampling model (cf. General Discussion). The depicted predictions are based on 1,000 replications with best-fitting parameter estimates $\alpha = 11.66$, $\lambda = 0.99$, and $\phi = 0.61$, and root mean square deviation = 0.32. 0-MI, no misinformation control condition; 1-MI, misinformation presented once; 3-MI, misinformation presented three times; 0-R, no retraction control conditions; 1-R, retraction presented once; 3-R, retraction presented three times



and three retractions reduced it further (contrasts 6 and 7), without, however, eliminating the continued influence of misinformation. Surprisingly, the effect of a single exposure of misinformation was reduced equally by one or three retractions; that is, in this case, three retractions failed to reduce the continued influence effect below the level achieved with one retraction (contrasts 1–3 and 9), and this level was significantly above that in the 0-MI control condition (contrast 0).

Excluding participants who did not acknowledge the retraction in the manipulation-check questions ($n = 12$, thus leaving between 18 and 21 participants per condition) did not change this pattern of results.

Recall Mean recall rates varied between .68 (1-MI/1-R) and .78 (3-MI/3-R) across the six experimental conditions.

Table 1 Contrasts calculated on inference scores in Experiment 1

Contrast	$t(154)$	p	Cohen's d
0 0-MI/0-R vs. 1-MI/3-R	3.06	<.01	.49
1 1-MI/0-R vs. 1-MI/1-R	5.50	<.001	.89
2 1-MI/0-R vs. 1-MI/3-R	5.41	<.001	.87
3 1-MI/1-R vs. 1-MI/3-R	<1		
4 1-MI/0-R vs. 3-MI/0-R	2.48	.01	.40
5 3-MI/0-R vs. 3-MI/1-R	5.09	<.001	.82
6 3-MI/0-R vs. 3-MI/3-R	8.25	<.001	1.33
7 3-MI/1-R vs. 3-MI/3-R	3.15	<.01	.51
8 1-MI/1-R vs. 3-MI/1-R	2.88	<.01	.46
9 1-MI/3-R vs. 3-MI/3-R	<1		

A two-way ANOVA returned no significant effects, $F_s < 1.5$, $p_s > .2$.²

Awareness of retraction Mean rates of acknowledgment across conditions ranged from .63 to .83. Although 3-R conditions yielded higher rates (.78) than 1-R conditions (.67), a two-way ANOVA yielded no significant effects, $F_s < 1.1$.

Discussion

Experiment 1 produced several noteworthy findings. First, in line with previous research, we found that even multiple retractions were insufficient to eliminate the continued influence of misinformation completely (cf. Bush, Johnson, & Seifert, 1994; Ecker et al., 2011).

Second, when misinformation was encoded once, there was a low but significant level of continued influence, and this influence was independent of the strength of retraction. In other words, after relatively weak encoding of misinformation, its influence was significant even if the retraction was strong. This corroborates research that has found it difficult to eliminate effects of misinformation, such as that of Ecker et al. (2010), who combined explicit warnings with the provision of a causal alternative but still found significant levels of continued influence after administering this combined manipulation.

² Visual data inspection and additional analyses of covariance ascertained that misinformation effects (in both experiments) were not mediated by recall performance.

Third, when misinformation was encoded three times but retracted only once, a relatively large continued influence effect was observed. Only repeated retractions were able to reduce this effect to the level elicited by one encoding of misinformation. The effectiveness of multiple retractions after strong encoding of misinformation does not support concerns that multiple retractions could enhance continued influence by increasing familiarity of the misinformation (Schwarz et al., 2007; Skurnik et al., 2005; see also Hintzman, 2010). It follows that the so-called backfire effects of retractions (Nyhan & Reifler, 2010; Pickel, 1995) may apply primarily to areas such as political beliefs or judicial settings, in which preexisting attitudes play a more important role for behavior.

The fact that there were no significant differences between conditions in fact recall and awareness of the retraction suggests that the differential pattern of the continued influence effect cannot be attributed to differences in overall memory strength.

Experiment 2

In Experiment 2, we implemented a different strength manipulation by introducing a cognitive load. It is well established that cognitive load—that is, the division of attention between two tasks—can have debilitating effects on memory retrieval (e.g., Craik, Naveh-Benjamin, Govoni, & Anderson, 1996) and can reduce depth of encoding and impede strategic processes (Magliano & Radvansky, 2001). Cognitive load at misinformation encoding should therefore reduce the continued influence of misinformation.

In contrast, cognitive load during retraction encoding should enhance the continued influence effect inasmuch as load impairs the updating of the situation model. Preliminary support for this idea was provided by Gilbert, Krull, and Malone (Gilbert et al. 1990), who found that imposing a cognitive load during immediate retraction of a proposition (of the type “an X is a Y”) increased the likelihood that retracted (and hence false) propositions would later be considered true.

Experiment 2 again involved the warehouse fire scenario. Cognitive load was imposed either when the misinformation was encoded or at the stage of retraction. The design again involved two between-subjects factors: load at misinformation (load vs. no load; conditions L-MI and noL-MI) and retraction (no retraction, load at retraction, no load at retraction; conditions 0-R, L-R, and noL-R, respectively). As in Experiment 1, the no-retraction conditions served as ceilings to assess the effects of retraction. The no-misinformation control group of Experiment 1 was dropped, given the negligible level of misinformation references in Experiment 1.

Method

Participants A sample of 138 undergraduates (95 females) participated, and participants were randomly assigned to conditions ($N = 23$ per condition).

Stimuli and procedure The script was similar to that used in Experiment 1 but consisted of only 14 messages, with the misinformation contained in Message 5 and the retraction in Message 10. After reading the script, participants additionally recalled a summary of the study (following precedents; Johnson & Seifert, 1994; Wilkes & Leatherbarrow, 1988). The questionnaire was identical to the one used in Experiment 1, and questions were given in the same order.

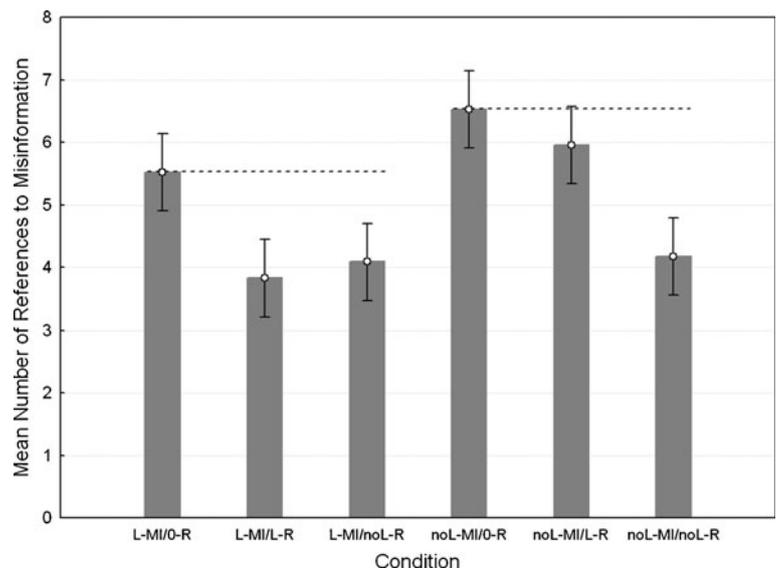
In all conditions, seven digits were presented in between each pair of messages. Participants were given 4 sec to either read aloud (no load) or read aloud and memorize (during load) these digits. After the subsequent message, participants either recalled and wrote down the memorized digits (during load) or read aloud the same re-presented digits (no load). The memory load was imposed during messages 4–9 in L-MI conditions, and during messages 9–14 in L-R conditions, thus bracketing the crucial misinformation and retraction messages, respectively. Serial recall-in-position accuracy in the load conditions was good ($M = .49$; $SE = .01$) and clearly above chance, $t(91) = 27.30$, $p < .001$, which demonstrates the effectiveness of our load manipulation. There were no significant differences in digit recall across conditions, $F(3, 88) = 2.26$, $p = .09$.

Results

Coding procedure The data were scored as in Experiment 1, except that references to misinformation made during free recall were also counted (Johnson & Seifert, 1994; Wilkes & Leatherbarrow, 1988). We expected numerically higher levels of references to misinformation (as compared with those in Experiment 1), which had the advantage of avoiding floor effects potentially hampering the assessment of retraction effects.

Inferences Figure 2 displays the mean inference scores. A two-way ANOVA yielded significant main effects of load at misinformation, $F(1, 132) = 4.54$, $p = .04$, $\eta^2 = .03$, and retraction, $F(2, 132) = 4.77$, $p = .01$, $\eta^2 = .07$. Hence, overall, misinformation effects were higher when misinformation was encoded without load, and retractions reduced misinformation effects. Planned contrasts (cf. Table 2), however, demonstrated that with no load at misinformation encoding, only a retraction likewise encoded without load significantly reduced misinformation effects (contrast 6),

Fig. 2 Mean number of references to misinformation for all conditions in Experiment 2; error bars show standard errors of the mean. *L-MI*, misinformation encoded under cognitive load; *noL-MI*, no cognitive load at misinformation encoding; *0-R*, no retraction control conditions; *L-R*, retraction encoded under cognitive load; *noL-R*, no cognitive load at retraction encoding



whereas a retraction under load was ineffective (contrast 5; see also contrast 7). When misinformation was encoded under load, both types of retraction (with or without load) were equally effective in reducing misinformation effects (contrasts 1–3), which nevertheless remained quite substantially above zero.

Excluding participants who did not acknowledge the correction in the manipulation-check questions from the retraction conditions ($n = 35$, leaving 13–16 participants in each condition) did not change the overall pattern of effects.

Recall Fact recall rates across conditions varied from .37 (noL-MI/L-R) to .50 (noL-MI/noL-R). This numerical difference fell short of significance, $F(5, 122) = 1.63$, $p = .16$.³

Awareness of retraction Acknowledgment of retraction scores varied from .35 (noL-MI/L-R) to .48 (noL-MI/noL-R). This difference also fell short of significance, $F < 1$.

Discussion

The aim of Experiment 2 was to investigate the effects of a cognitive load, imposed either during the encoding of misinformation or during its retraction, on the continued influence effect. Misinformation had larger effects if encoded without load, and continued influence was reduced only if the retraction was encoded with full attentional resources (without load). As in Experiment 1, the effects of relatively weakly encoded misinformation (under load)

were reduced by both weak (load) and strong (no load) retractions to the same degree. Hence, the outcome perfectly mimicked the results of Experiment 1 in that after relatively weak encoding of misinformation, even a comparatively stronger retraction failed to reduce continued influence below the level achieved by a weaker retraction. This is further evidence that misinformation effects are very difficult to reduce below a certain level, be it by strengthening the retraction or by other means, such as the provision of causal alternatives and explicit warnings (cf. Ecker et al., 2010).

One open question is why, when misinformation was encoded without load, the effectiveness of a retraction was so drastically reduced under cognitive load? The failure of a retraction is in line with previous research (Johnson & Seifert, 1994), but it is at odds with the results of Experiment 1. At the moment, it remains unclear why retractions sometimes reduce the continued influence effect but at other times fail to do so; however, the finding that

Table 2 Contrasts calculated on inference scores in Experiment 2

Contrast	$t(132)$	p	Cohen's d
1 L-MI/0-R vs. L-MI/L-R	1.95	.05	.34
2 L-MI/0-R vs. L-MI/noL-R	1.65	.10	.29
3 L-MI/L-R vs. L-MI/noL-R	<1		
4 L-MI/0-R vs. noL-MI/0-R	1.15	.25	.20
5 noL-MI/0-R vs. noL-MI/L-R	<1		
6 noL-MI/0-R vs. noL-MI/noL-R	2.69	<.01	.47
7 noL-MI/L-R vs. noL-MI/noL-R	2.05	.04	.36
8 L-MI/L-R vs. noL-MI/L-R	2.45	.02	.43
9 L-MI/noL-R vs. noL-MI/noL-R	<1		

³ Due to a logistical problem, fact recall data was available for only 13 subjects of the noL-MI/0-R conditions.

retractions never eliminate continued influence altogether is pervasive and robust.

General discussion

In two experiments, we manipulated the strength of misinformation encoding and its retraction. We found that stronger encoding of misinformation resulted in increased levels of continued influence. This is not unexpected, because repeated encoding typically enhances memory (e.g., Verkoeijen et al., 2004). However, this result contradicts van Oostendorp's (1996) suggestion that a strong representation of misinformation could facilitate its updating.

The results of both experiments also suggested that greater misinformation effects required stronger retractions to substantially reduce continued influence. More interestingly, however, the results of both experiments suggested that the strength of retraction is immaterial if misinformation is only encoded relatively weakly. Although this replicates previous research (van Oostendorp & Bonebakker, 1999), the pattern remains to be explained.

One way to interpret this finding, in terms of the dual-process mechanism discussed earlier, is that strategic processing aiming to minimize illogical inferences based on misinformation can only counteract the automatic activation of misinformation to the degree that the person is in fact aware of the automatic influence. Wilson and Brekke (1994) argued that unintentional effects of inappropriate information mainly occur because people are unaware of the extent of these influences. In fact, there are instances reported in the literature in which participants failed to avoid influences of automatic processing in their judgments despite efforts to minimize them, but these were implicit effects (e.g., a larger weapons-false-alarm effect when primed with a Black vs. a White face; Payne, Lambert, & Jacoby, 2002). Therefore, we consider an explanation along these lines speculative because continued influence is measured by direct, explicit inferences.

We thus prefer to seek an explanation by modeling the detailed underlying memorial processes. Why do people remember the retraction but nevertheless use misinformation in their reasoning? And, in particular, why does it seem so hard to reduce misinformation effects below a certain level? We focused our attention on the junction of memory and reasoning, in particular the way in which the memory system might support inferences. The simplest mechanism would involve random sampling: If misinformation were randomly sampled from memory, and hence were more likely to be sampled if more misinformation was represented, and if the impact of misinformation were largely but not entirely offset by retractions, could this explain the observed pattern?

We fleshed out this potential explanation by designing a simple sampling model that relied on the following assumptions:

1. Inferences depend on drawing a limited set of samples ($N_s = 4$ in our instantiation) from an ensemble of facts in memory.
2. The memorial ensemble ($N_m = 12$) holds pieces of event information, including misinformation.
3. Repeated (or strong) encoding of a specific piece of information results in the creation of multiple tokens.
4. The strength of repeated tokens of the same piece of information declines according to $s = \alpha \exp(-\lambda i)$, where i runs from 0 to one less than the number of repetitions, λ is a parameter determining the rate of decline over repetitions, and α is an arbitrary scaling parameter. This equation instantiates “novelty-sensitive” encoding, as embodied in numerous memory models (in particular SOB; Farrell & Lewandowsky, 2002; Lewandowsky & Farrell, 2008). There is considerable evidence for the notion that repetitions are encoded with less strength than in the first presentation (e.g., Oberauer & Lewandowsky, 2008).
5. Retractions are represented as negation tags linked to specific tokens (cf. Gilbert et al., 1990), which implies that retractions can be effective only if the associated misinformation token is sampled.
6. If a retraction tag is presented, it is assumed to have been encoded, and it reduces—but does not eliminate—the impact of the associated piece of misinformation by a factor ϕ ($\phi < 1$).
7. Crucially, each misinformation token can be offset only by one negation tag.

This model reproduces the data from Experiment 1, as shown in Fig. 1. Predictions are obtained by summing the final strength values of all misinformation tokens, thus using the scaling parameter α to convert memory strength into number of inferences. (The MATLAB code can be accessed via an online supplement at <http://www.cogsciwa.com>).

This simple sampling mechanism gives rise to a pattern in which the impact of misinformation, once in the cognitive system, is difficult to drive below a certain level of “irreducible persistence,” because a retraction can be coupled with a misinformation token only once (assumption g, above). Hence, unless multiple misinformation tokens are present, multiple retractions will be no more effective than a single one. To conclude, this is the first computational model to be applied to the continued influence effect of misinformation. Of course, at this stage, this is only an illustration of how an exemplar-based sampling model might account for the pervasive finding that continued influence is extremely difficult to eliminate (cf. Ecker et al., 2011; Ecker et al., 2010; van Oostendorp

& Bonebakker, 1999); more empirical work is needed to test the crucial model assumptions outlined previously in this article.

Practical implications

The practical implications of the present research are clear: If misinformation is encoded strongly, the level of continued influence will significantly increase, *unless* the misinformation is also retracted strongly. Hence, if information that has had a lot of news coverage is found to be incorrect, the retraction will need to be circulated with equal vigor, or else continued influence will persist at high levels. Of course, in reality, initial reports of an event, which may include misinformation (e.g., that a person of interest has committed a crime or that a country seeks to hide WMDs), may attract more interest than their retraction. Moreover, retractions apparently need full attentional resources to become effective; hence, retractions processed during conditions of divided attention (e.g., when listening to the news while driving a car) may remain ineffective.

Acknowledgements Preparation of this manuscript was facilitated by a Discovery Grant and an Australian Professorial Fellowship from the Australian Research Council to S. L. The lab's website is located at <http://www.cogsciwa.com>. We thank Charles Hanich for research assistance.

References

- Allport, F. H., & Lepkin, M. (1945). Wartime rumors of waste and special privilege: Why some people believe them. *Journal of Abnormal and Social Psychology*, *40*, 3–36.
- 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.
- Baron-Cohen, S. (2009). Media distortion damages both science and journalism. *New Scientist*, *2701*, 26–28.
- Bush, J. G., Johnson, H. M., & Seifert, C. M. (1994). The implications of corrections: Then why did you mention it? In A. Ram & K. Eiselt (Eds.), *Proceedings of the sixteenth annual conference of the cognitive science society* (pp. 112–117). Hillsdale, NJ: Erlbaum.
- Chabot, R. J., Miller, T. J., & Juola, J. F. (1976). The relationship between repetition and depth of processing. *Memory & Cognition*, *4*, 677–682.
- Craik, F. I. M., Naveh-Benjamin, M., Govoni, R., & Anderson, N. D. (1996). The effects of divided attention on encoding and retrieval processes in human memory. *Journal of Experimental Psychology: General*, *125*, 159–180.
- Eakin, D. K., Schreiber, T. A., & Sergent-Marshall, S. (2003). Misinformation effects in eyewitness memory: The presence and absence of memory impairment as a function of warning and misinformation accessibility. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 813–825.
- Ecker, U. K. H., Lewandowsky, S., & Apai, J. (2011). 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., Oberauer, K., & Chee, A. E. H. (2010). The components of working memory updating: An experimental decomposition and individual differences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *36*, 170–189.
- Ecker, U. K. H., Lewandowsky, S., & Tang, D. T. W. (2010). Explicit warnings reduce but do not eliminate the continued influence of misinformation. *Memory & Cognition*, *38*, 1087–1100.
- Farrell, S., & Lewandowsky, S. (2002). An endogenous distributed model of ordering in serial recall. *Psychonomic Bulletin & Review*, *9*, 59–79.
- Gilbert, D. T., Krull, D. S., & Malone, P. S. (1990). Unbelieving the unbelievable: Some problems in the rejection of false information. *Journal of Personality and Social Psychology*, *59*, 601–613.
- Hintzman, D. L. (2010). How does repetition affect memory? Evidence from judgments of recency. *Memory & Cognition*, *38*, 102–115.
- Hupbach, A., Gomez, R., Hardt, O., & Nadel, L. (2007). Reconsolidation of episodic memories: A subtle reminder triggers integration of new information. *Learning & Memory*, *14*, 47–53.
- Jacoby, L. L. (1999). Ironic effects of repetition: Measuring age-related differences in memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 3–22.
- 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, and Cognition*, *20*, 1420–1436.
- Kull, S., Ramsay, C., & Lewis, E. (2003). Misperceptions, the media, and the Iraq war. *Political Science Quarterly*, *118*, 569–598.
- Lewandowsky, S., & Farrell, S. (2008). Short-term memory: New data and a model. *Psychology of Learning and Motivation*, *49*, 1–48.
- Lewandowsky, S., Stritzke, W. G. K., Oberauer, K., & Morales, M. (2005). Memory for fact, fiction and misinformation: The Iraq War 2003. *Psychological Science*, *16*, 190–195.
- Lewandowsky, S., Stritzke, W., Oberauer, K., & Morales, M. (2009). Misinformation and the ‘War on Terror’: When memory turns fiction into fact. In W. Stritzke, S. Lewandowsky, D. Denemark, J. Clare, & F. Morgan (Eds.), *Terrorism and torture: An interdisciplinary perspective* (pp. 179–203). Cambridge, UK: Cambridge University Press.
- Lewis, C., & Reading-Smith, M. (2008). Iraq: The war card. *The Center for Public Integrity*. Retrieved December 23, 2009, from <http://projects.publicintegrity.org/WarCard/>
- Magliano, J. P., & Radvansky, G. A. (2001). Goal coordination in narrative comprehension. *Psychonomic Bulletin & Review*, *8*, 372–376.
- Nyhan, B., & Reifler, J. (2010). When corrections fail: The persistence of political misperceptions. *Political Behavior*, *32*, 303–330.
- Oberauer, K., & Lewandowsky, S. (2008). Forgetting in immediate serial recall: Decay, temporal distinctiveness, or interference? *Psychological Review*, *115*, 544–576.
- Payne, B. K., Lambert, A. J., & Jacoby, L. L. (2002). Best laid plans: Effects of goals on accessibility bias and cognitive control in race-based misperceptions of weapons. *Journal of Experimental Social Psychology*, *38*, 384–396.
- Pickel, K. L. (1995). Inducing jurors to disregard inadmissible evidence: A legal explanation does not help. *Law and Human Behavior*, *19*, 407–424.
- Radvansky, G. A., & Copeland, D. E. (2001). Working memory and situation model updating. *Memory & Cognition*, *29*, 1073–1080.
- Schul, Y., & Mazursky, D. (1990). Conditions facilitating successful discounting in consumer decision making. *Journal of Consumer Research*, *15*, 442–451.

- Schwarz, N., Sanna, L. J., Skurnik, I., & Yoon, C. (2007). Metacognitive experiences and the intricacies of setting people straight: Implications for debiasing and public information campaigns. *Advances in Experimental Social Psychology*, 39, 127–161.
- Sewell, D. K., & Lewandowsky, S. (2011). Restructuring partitioned knowledge: The role of recoordination in category learning. *Cognitive Psychology*, 62, 81–122.
- Skurnik, I., Yoon, C., Park, D. C., & Schwarz, N. (2005). How warnings about false claims become recommendations. *Journal of Consumer Research*, 31, 713–724.
- van Oostendorp, H. (1996). Updating situation models derived from newspaper articles. *Medienpsychologie*, 8, 21–33.
- van Oostendorp, H., & Bonebakker, C. (1999). Difficulties in updating mental representations during reading news reports. In H. van Oostendorp & S. R. Goldman (Eds.), *The construction of mental representations during reading* (pp. 319–339). Mahwah, NJ: Erlbaum.
- Verkoeijen, P. P. J. L., Rikers, R. M. J. P., & Schmidt, H. G. (2004). Detrimental influence of contextual change on spacing effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 796–800.
- 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.
- Wilson, T. D., & Brekke, N. (1994). Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin*, 116, 117–142.