

PSYCHONOMIC BULLETIN & REVIEW

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

Journal:	<i>Psychonomic Bulletin & Review</i>
Manuscript ID:	PBR-BR-10-056.R3
Manuscript Type:	Brief Report
Date Submitted by the Author:	n/a
Complete List of Authors:	Ecker, Ullrich; University of Western Australia, Psychology Lewandowsky, Stephan; University of Western Australia, Psychology Swire, Briony; University of Western Australia, Psychology Chang, Darren; University of Western Australia, Psychology
Keywords:	episodic memory and recall, text comprehension, reasoning, misinformation

SCHOLARONE™
Manuscripts

Running head: ENCODING AND RETRACTION OF MISINFORMATION

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

Ullrich K. H. Ecker, Stephan Lewandowsky, Briony Swire, and Darren Chang

School of Psychology

University of Western Australia

Ullrich K. H. Ecker

School of Psychology

University of Western Australia

Crawley, W.A. 6009, Australia

Telephone: +618 6488 3266

Fax: +618 6488 1006

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

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. The present study 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.

1
2
3
4 Correcting false information in memory: Manipulating the strength of misinformation
5
6 encoding and its retraction
7

8 Once encoded, information may continue to influence reasoning even if it later turns
9
10 out to be incorrect. The persistent reliance on such misinformation, even when people can
11
12 recall a correction or retraction, has been labeled the continued influence effect (H. Johnson &
13
14 Seifert, 1994). For example, if a fictional character is accused of a crime but is later
15
16 exonerated, people continue to use the outdated misinformation (that the person is guilty) in
17
18 subsequent inferences even if they recall the correction. The continued influence effect has
19
20 been demonstrated in many settings (e.g., Ecker, Lewandowsky, & Apai, in press; Ecker,
21
22 Lewandowsky, & Tang, 2010; H. Johnson & Seifert, 1994; van Oostendorp, 1996; Wilkes &
23
24 Leatherbarrow, 1988). In the real world, continued belief in unsubstantiated claims can have
25
26 serious implications, as in the case of the purported link between certain vaccines and autism
27
28 (Baron-Cohen, 2009) or between Iraq and weapons of mass destruction (WMDs; Kull,
29
30 Ramsay & Lewis, 2003; Lewandowsky, Stritzke, Oberauer, & Morales, 2005, 2009).
31
32
33
34
35

36 Two different approaches have been put forward to explain the continued influence
37
38 effect. One of these refers to “mental event models” (H. Johnson & Seifert, 1994; Wilkes &
39
40 Leatherbarrow, 1988). People are thought to build mental models of unfolding events but
41
42 seem reluctant to dismiss key information in their model (e.g., what caused an event) when no
43
44 plausible alternative exists to fill the void. Accordingly, the provision of alternative causal
45
46 information (e.g., presentation of an alternative suspect) has long been the only factor known
47
48 to reduce the continued influence effect (H. Johnson & Seifert, 1994). When no alternative is
49
50 presented, people prefer an inconsistent event model to an incomplete event model. Hence, in
51
52 their inferential reasoning they may rely on outdated information despite knowing that it is
53
54 false, rather than acknowledging the lack of valid information available.
55
56
57
58
59
60

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52

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 (1998) 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 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.

30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

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 two years following 9/11 (Lewis & Reading-Smith, 2008). It is possible that the reiteration of this misinformation (e.g., 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.

53
54
55
56
57
58
59
60

According to the event model approach, the initial integration of information into the event model is more readily performed than 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

1
2
3 event model, and the new information (e.g., *unclear what caused Y but it was 'not X'*) must be
4
5 integrated (cf. Ecker, Lewandowsky, Oberauer, & Chee, 2010; Radvansky & Copeland, 2001).
6

7
8 This suggests that the retraction may profit more from repetition than the encoding of
9
10 misinformation and hence we may expect differential effects of strengthening.
11

12
13 In contrast, the dual-processing account would predict the opposite. Eakin, Schreiber,
14
15 and Sergent-Marshall (2003) demonstrated that misinformation effects could be suppressed
16
17 by explicit warnings, which foster strategic monitoring processes (cf. Ecker, Lewandowsky,
18
19 & Tang, 2010). This suggests that the effects of retractions are mainly carried by strategic
20
21 processing. However, warnings failed to reduce misinformation effects when misinformation
22
23 was presented repeatedly, which according to Eakin et al. mainly fostered automatic retrieval.
24
25 On the assumption that strengthening information encoding will mainly affect automatic
26
27 processes, repetition may have a greater impact on the encoding of misinformation than its
28
29 retraction.
30
31
32

33
34 The present study is the first to systematically investigate how the strength with which
35
36 misinformation is encoded, and the vigor with which it is later retracted, affects the continued
37
38 influence effect. It is well documented that repetition enhances belief in the truth of repeated
39
40 assertions (e.g., Allport & Lepkin, 1945; Weaver, Garcia, Schwarz, & Miller, 2007) as well as
41
42 memory more generally. This is especially true if repetition occurs with some temporal
43
44 spacing or in different contexts (cf. Chabot, Miller, & Juola, 1976; Verkoeijen, Rikers, &
45
46 Schmidt, 2004). This in turn suggests that repetition unfolds its effects by associating
47
48 information with various contexts or sources, which could serve as retrieval cues and/or lend
49
50 more credibility to a repeatedly encoded piece of information. Hence, enhanced encoding
51
52 could increase the continued influence effect because repeated misinformation may become
53
54 harder to retract (cf. Schul & Mazursky, 1990).
55
56
57
58
59
60

1
2
3
4 Alternatively, however, it has been suggested that enhanced encoding of
5
6 misinformation may *reduce* its continued influence, as memory updating may be more
7
8 efficient when the initial information—despite being false—is well represented and active in
9
10 memory (van Oostendorp, 1996). The idea that only something that is well represented in
11
12 memory can be easily updated is in line with at least two related areas of enquiry. First,
13
14 reconsolidation theory claims that information needs to be activated in order to be updated
15
16 (e.g., Hupbach, Gomez, Hardt, & Nadel, 2007). Second, in the categorization and problem
17
18 solving literature it has become clear that knowing something well is no barrier for, or may
19
20 even benefit, knowledge restructuring (i.e., shift to an alternate strategy; Sewell &
21
22 Lewandowsky, 2011). Finally, in terms of dual processes, repetition will not only enhance
23
24 automatic retrieval (as suggested by Eakin et al., 2003), but will usually lead to improved
25
26 controlled memory processes as well, in particular improved source memory (e.g., Jacoby,
27
28 1999). Hence, inasmuch as factors such as source confusion are reduced by repetition,
29
30 continued influence could also be reduced by repetition of misinformation.
31
32
33
34
35

36 Concerning the retraction, intuitively, if a statement is retracted with greater emphasis,
37
38 one might expect less continued influence. However, the only study known to us that
39
40 manipulated the strength of the retraction (van Oostendorp & Bonebakker, 1999) found
41
42 continued influence to be unaffected by repetition of a retraction—two retractions were found
43
44 to be as (in-)effective as one. Moreover, the literature on metacognitive effects of repetition
45
46 has shown that ironically, because misinformation is often repeated when it is retracted, more
47
48 frequent retraction of misinformation can paradoxically enhance its impact even after
49
50 relatively short retention intervals of 30 minutes (Schwarz et al., 2007; Skurnik, Yoon, Park,
51
52 & Schwarz, 2005). In other words, the retraction could serve as a recursive reminder of the
53
54 misinformation (Hintzman, 2010). Such backfire effects of retractions have been observed in
55
56
57
58
59
60

1
2
3 examinations of the effects of retractions on political misperceptions (Nyhan & Reifler, 2010)
4
5 and mock juror behavior (Pickel, 1995), and are obviously a reason for concern.
6
7

8 In summary, it is unclear whether continued-influence effects are necessarily enhanced
9
10 and reduced, respectively, by repetition or other means of strengthening the initial encoding of
11
12 misinformation or its retraction.
13

14
15 We present two experiments that manipulated the strength with which misinformation
16
17 was encoded and retracted. Experiment 1 orthogonally varied the number of repetitions of the
18
19 misinformation and its retraction; Experiment 2 used a cognitive load manipulation at the
20
21 encoding or the retraction stage. In both experiments, participants received an adaptation of a
22
23 much-used warehouse fire script (H. Johnson & Seifert, 1994; Wilkes & Leatherbarrow,
24
25 1988), in which a fire was initially reported to have been caused by volatile materials stored
26
27 negligently in a closet, with subsequent reports retracting this, stating that the closet had been
28
29 empty.
30
31
32

33 34 Experiment 1

35
36 Two between-subjects factors were fully crossed. Factors were the strength-of-
37
38 misinformation (1 or 3 repetitions), and the strength-of-retraction (0, 1, or 3 repetitions).
39
40 Additionally, a control group with no mention of volatile materials was tested.
41
42

43 *Method*

44
45
46 *Participants.* One hundred sixty-one undergraduates from the University of Western
47
48 Australia (108 females) participated and were randomly assigned to conditions¹ ($N = 23$ per
49
50 condition).
51

52
53 *Stimuli.* Participants received 17 messages, each printed on a separate page. The *0-MI*
54
55 control condition featured no statements referring to volatile materials, and obviously no
56
57 retraction either. In the *1-* and *3-MI* conditions, a statement regarding the presence of volatile
58
59 materials appeared once (Message 6) or three times (Messages 6, 7, and 8), respectively. In
60

1
2
3 the 3-R conditions, the retraction was repeated three times (Messages 13, 14, and 15) as
4
5 opposed to only once in the 1-R conditions (Message 14), and no retraction at all in the 0-R
6
7 conditions. Repetitions were presented in different contexts (e.g., as a radio transmission from
8
9 a police investigator / as information passed on to the fire captain / in a public radio
10
11 announcement). This was done to make the repetitions appear more natural and to enhance
12
13 their potential impact by increasing contextual variation (e.g., Verkoeijen et al., 2004). Scripts
14
15 were of equal length in all conditions; filler messages were added where needed.
16
17
18

19
20 *Procedure.* Participants read the messages aloud at their own pace, without
21
22 backtracking. After an unrelated 10-minute distractor task, participants received an open-
23
24 ended questionnaire, consisting of 10 causal inference questions (e.g., *What could have*
25
26 *caused the explosions?*), 10 fact questions (e.g., *What time was the fire eventually put out?*),
27
28 and 2 manipulation-check questions, targeting awareness of the retraction (e.g., *Was any of*
29
30 *the information in the story subsequently corrected or altered? And if so, what was it?*),
31
32 always administered in this order.
33
34
35

36 37 *Results*

38
39 Analysis focused on three dependent measures: the number of references to
40
41 misinformation (i.e., the inference score), the accuracy of recall, and acknowledgment of the
42
43 retraction. References to misinformation (viz. negligently stored volatile materials) were only
44
45 counted if they were causal and uncontroverted.
46
47

48
49 *Coding procedure.* Responses were tallied by a naïve scorer following a scoring guide.
50
51 Inter-rater reliability with a second scorer was high ($r = .97/.82/.94$ for fact-recall, inference,
52
53 and manipulation-check scores, respectively, based on a sample of 18 questionnaires).
54

55
56 *Inferences.* Mean inference scores are shown in Figure 1. The 0-MI and 0-R conditions
57
58 provide empirical baselines for interpretation of the remaining experimental conditions and
59
60 are represented by the dotted lines. The 0-MI control condition expectedly yielded few

1
2
3 spontaneous references to misinformation (significantly fewer than the *1-MI/3-R* condition
4
5 which was expected to have the lowest score of the remaining conditions; all contrasts are
6
7 given in Table 1).
8
9

10 A two-way ANOVA on the six experimental conditions yielded significant main effects
11
12 of strength-of-misinformation, $F(1,132) = 7.32, p < .01, \eta^2 = .05$, and strength-of-retraction,
13
14 $F(2,132) = 45.02, p < .001, \eta^2 = .41$, which were qualified by a marginally significant
15
16 interaction, $F(2,132) = 2.74, p = .07, \eta^2 = .04$. Planned contrasts (cf. Table 1) demonstrated
17
18 that, not surprisingly, that repeated misinformation encoding led to stronger misinformation
19
20 effects when there was no or only one retraction (contrasts 4 and 8). After three presentations
21
22 of misinformation, one retraction reduced reliance on misinformation (contrast 5), and three
23
24 retractions reduced it further (contrasts 6 and 7), without however eliminating the continued
25
26 influence of misinformation. Surprisingly, the effect of a single exposure of misinformation
27
28 was reduced equally by one or three retractions, that is, in this case three retractions failed to
29
30 reduce the continued influence effect below the level achieved with one retraction (contrasts
31
32 1-3, and 9), and this level was significantly above the *0-MI* control condition (contrast 0).
33
34
35
36
37
38

39 Excluding participants who did not acknowledge the retraction in the manipulation-
40
41 check questions ($n = 12$, thus leaving between 18 and 21 participants per condition) did not
42
43 change this pattern of results.
44
45

46 *Recall.* Mean recall rates varied between .68 (*1-MI/1-R*) and .78 (*3-MI/3-R*) across the
47
48 six experimental conditions. A two-way ANOVA returned no significant effects, F 's < 1.5 ,
49
50 p 's $> .2$.²
51
52

53 *Awareness of retraction.* Mean rates of acknowledgment across conditions ranged
54
55 from .63 to .83. Although *3-R* conditions yielded higher rates (.78) than *1-R* conditions (.67),
56
57 a two-way ANOVA yielded no significant effects, F 's < 1.1 .
58
59
60

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, H. Johnson, & Seifert, 1994; Ecker et al., in press).

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 Ecker, Lewandowsky, and Tang (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.

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, where pre-existing 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

Experiment 2 implemented a different strength manipulation by introducing a cognitive load. It is well established that cognitive load, i.e., 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 (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 were randomly assigned to conditions ($N = 23$ per condition).

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Stimuli and Procedure. The script was similar to 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; H. 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 four seconds 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 (H. Johnson & Seifert, 1994; Wilkes & Leatherbarrow, 1988). We expected numerically higher levels of references to misinformation (as compared to 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

1
2
3 higher when misinformation was encoded without load, and retractions reduced
4
5 misinformation effects. Planned contrasts (cf. Table 2), however, demonstrated that with no
6
7 load at misinformation encoding, only a retraction likewise encoded without load significantly
8
9 reduced misinformation effects (contrast 6), whereas a retraction under load was ineffective
10
11 (contrast 5; see also contrast 7). When misinformation was encoded under load, both types of
12
13 retraction (with or without load) were equally effective in reducing misinformation effects
14
15 (contrasts 1-3), which nevertheless remained quite substantially above zero.
16
17

18
19
20 Excluding participants who did not acknowledge the correction in the manipulation-
21
22 check questions from the retraction conditions ($n = 35$, leaving 13-16 participants in each
23
24 condition) did not change the overall pattern of effects.
25
26

27
28 *Recall.* Fact recall rates across conditions varied from .37 (*noL-MI/L-R*) to .50 (*noL-*
29
30 *MI/noL-R*). This numerical difference fell short of significance, $F(5,122) = 1.63, p = .16$.³
31

32
33 *Awareness of retraction.* Acknowledgment of retraction scores varied from .35 (*noL-*
34
35 *MI/L-R*) to .48 (*noL-MI/noL-R*). This difference also fell short of significance, $F < 1$.
36

37 *Discussion*

38
39 The aim of Experiment 2 was to investigate the effects of a cognitive load, imposed
40
41 either during the encoding of misinformation or during its retraction, on the continued
42
43 influence effect. Misinformation had larger effects if encoded without load, and only if the
44
45 retraction was encoded with full attentional resources (without load) was continued influence
46
47 reduced. As in Experiment 1, the effects of relatively weakly encoded misinformation (under
48
49 load) were reduced by both weak (load) and strong (no load) retractions to the same degree.
50
51 Hence, the outcome perfectly mimicked the results of Experiment 1 in that after relatively
52
53 weak encoding of misinformation, even a comparatively stronger retraction failed to reduce
54
55 continued influence below the level achieved by a weaker retraction. This is further evidence
56
57 that misinformation effects are very difficult to reduce below a certain level, be it by
58
59
60

1
2
3 strengthening the retraction or by other means such as the provision of causal alternatives and
4
5 explicit warnings (cf. Ecker, Lewandowsky, & Tang, 2010).
6
7

8 One open question is why, when misinformation was encoded without load, the
9
10 effectiveness of a retraction was so drastically reduced under cognitive load. The failure of a
11
12 retraction is in line with previous research (H. Johnson & Seifert, 1994), but it is at odds with
13
14 Experiment 1. At the moment, it remains unclear why retractions sometimes reduce the
15
16 continued influence effect but at other times fail to do so; however, the finding that retractions
17
18 never eliminate continued influence altogether is pervasive and robust,
19
20
21

22 General Discussion

23
24 In two experiments, we manipulated the strength of misinformation encoding and its
25
26 retraction. We found that stronger encoding of misinformation resulted in increased levels of
27
28 continued influence. This is not unexpected because repeated encoding typically enhances
29
30 memory (e.g., Verkoeijen et al., 2004). However, this result contradicts van Oostendorp's
31
32 (1996) suggestion that a strong representation of misinformation could facilitate its updating.
33
34
35

36 Both experiments also suggested that greater misinformation effects required stronger
37
38 retractions to substantially reduce continued influence. More interestingly, however, both
39
40 experiments suggested that the strength of retraction is immaterial if misinformation is only
41
42 encoded relatively weakly. Although this replicates previous research (van Oostendorp &
43
44 Bonebakker, 1999), the pattern remains to be explained.
45
46
47

48 One way to interpret this, in terms of the dual process mechanism discussed earlier, is
49
50 that strategic processing aiming to minimize illogical inferences based on misinformation can
51
52 only counteract the automatic activation of misinformation to the degree that the person is in
53
54 fact aware of the automatic influence. Wilson and Brekke (1994) argued that unintentional
55
56 effects of inappropriate information mainly occur because people are unaware of the extent of
57
58 these influences. In fact, there are instances reported in the literature where participants failed
59
60

1
2
3 to avoid influences of automatic processing in their judgments despite efforts to minimize
4 them, but these were implicit effects (e.g., a larger weapons-false-alarm effect when primed
5 with a Black vs. White face; Payne, Lambert, & Jacoby, 2002). Therefore, we consider an
6 explanation along these lines speculative because continued influence is measured by direct,
7 explicit inferences.
8
9

10
11 We thus prefer to seek an explanation by modeling the detailed underlying memorial
12 processes. Why do people remember the retraction but nevertheless use misinformation in
13 their reasoning? And in particular, why does it seem so hard to reduce misinformation effects
14 below a certain level? We focused our attention on the junction of memory and reasoning, in
15 particular the way in which the memory system might support inferences. The simplest
16 mechanism would involve random sampling: If misinformation were randomly sampled from
17 memory, and hence more likely to be sampled if more misinformation was represented, and if
18 the impact of misinformation was largely but not entirely offset by retractions, could this
19 explain the observed pattern?
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34

35
36 We fleshed out this potential explanation by designing a simple sampling model that
37 relied on the following assumptions: (a) Inferences depend on drawing a limited set of
38 samples ($N_s = 4$ in our instantiation) from an ensemble of facts in memory. (b) The memorial
39 ensemble ($N_m = 12$) holds pieces of event information, including misinformation. (c) Repeated
40 (or strong) encoding of a specific piece of information results in the creation of multiple
41 tokens. (d) The strength of repeated tokens of the same piece of information declines
42 according to $s = \alpha \exp(-\lambda i)$, where i runs from 0 to one less than the number of repetitions, λ
43 is a parameter determining the rate of decline over repetitions, and α is an arbitrary scaling
44 parameter. This equation instantiates “novelty-sensitive” encoding, as embodied in numerous
45 memory models (in particular SOB; Farrell & Lewandowsky, 2002; Lewandowsky & Farrell,
46 2008). There is considerable evidence for the notion that repetitions are encoded with less
47
48
49
50
51
52
53
54
55
56
57
58
59
60

1
2
3 strength than the first presentation (e.g., Oberauer & Lewandowsky, 2008). (e) Retractions are
4 represented as negation-tags linked to specific tokens (cf. Gilbert et al., 1990), which implies
5 that retractions can only be effective if the associated misinformation token is sampled. (f) If a
6 retraction tag is presented, it is assumed to have been encoded and it reduces—but does not
7 eliminate—the impact of the associated piece of misinformation by a factor ϕ ($\phi < 1$).
8
9

10
11
12
13 (g) Crucially, each misinformation token can only be offset by one negation-tag.
14

15
16
17 This model reproduces the data from Experiment 1, as shown in Figure 1. Predictions
18 are obtained by summing the final strength values of all misinformation tokens, thus using the
19 scaling parameter α to convert memory strength into number of inferences. (The MatLab code
20 can be accessed via an online supplement at...).

21
22
23 This simple sampling mechanism gives rise to a pattern where the impact of
24 misinformation, once in the cognitive system, is difficult to drive below a certain level of
25 “irreducible persistence,” because a retraction can only be coupled with a misinformation
26 token once (assumption g above). Hence, unless multiple misinformation tokens are present,
27 multiple retractions will be no more effective than a single one. To conclude, this is the first
28 computational model to be applied to the continued influence effect of misinformation. Of
29 course, at this stage this is only an illustration of how an exemplar-based sampling model
30 might account for the pervasive finding that continued influence is extremely difficult to
31 eliminate (cf. Ecker et al., in press; Ecker, Lewandowsky, & Tang, 2010; van Oostendorp &
32 Bonebakker, 1999); more empirical work is needed to test the crucial model assumptions
33 outlined above.
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52

53 Practical Implications

54
55 The practical implications of this research are clear: If misinformation is encoded
56 strongly, the level of continued influence will significantly increase, *unless* the
57 misinformation is also retracted strongly. Hence, if information that has had a lot of news
58
59
60

1
2
3 coverage is found to be incorrect, the retraction will need to be circulated with equal vigor, or
4
5 else continued influence will persist at high levels. Of course, in reality, initial reports of an
6
7 event, which may include misinformation (e.g., that a person of interest has committed a
8
9 crime or that a country seeks to hide WMDs), may attract more interest than their retraction.
10
11 Moreover, retractions apparently need full attentional resources to become effective, and
12
13 hence retractions processed during conditions of divided attention (e.g., when listening to the
14
15 news while driving a car) may remain ineffective.
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

For Review Only

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, March 28). 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: Lawrence Erlbaum Associates.

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. (in press). Terrorists brought down the plane!—No, actually it was a technical fault: Processing corrections of emotive information. *Quarterly Journal of Experimental Psychology*.

1
2
3 Ecker, U. K. H., Lewandowsky, S., Oberauer, K., & Chee, A. E. H. (2010). The
4 components of working memory updating: An experimental decomposition and individual
5 differences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36, 170-
6 189.
7

8
9
10
11
12 Ecker, U. K. H., Lewandowsky, S., & Tang, D. T. W. (2010). Explicit warnings
13 reduce but do not eliminate the continued influence of misinformation. *Memory & Cognition*,
14 38, 1087-1100.
15

16
17
18
19
20 Farrell, S., & Lewandowsky, S. (2002). An endogenous distributed model of ordering
21 in serial recall. *Psychonomic Bulletin & Review*, 9, 59-79.
22

23
24
25
26
27
28
29
30
31 Gilbert, D. T., Krull, D. S., & Malone, P. S. (1990). Unbelieving the unbelievable:
32 some problems in the rejection of false information. *Journal of Personality and Social*
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

48
49
50
51
52
53
54
55
56
57
58
59
60

55
56
57
58
59
60

1
2
3 Lewandowsky, S., & Farrell, S. (2008). Short-term memory: New data and a model.
4
5 *The Psychology of Learning and Motivation*, 49, 1-48.

6
7
8 Lewandowsky, S., Stritzke, W. G. K., Oberauer, K., & Morales, M. (2005). Memory
9
10 for fact, fiction and misinformation: The Iraq War 2003. *Psychological Science*, 16, 190-195.

11
12 Lewandowsky, S., Stritzke, W., Oberauer, K., & Morales, M. (2009). Misinformation
13
14 and the 'War on Terror': When memory turns fiction into fact. In W. Stritzke, S.
15
16 Lewandowsky, D. Denemark, J. Clare, & F. Morgan (Eds.), *Terrorism and torture: An*
17
18 *interdisciplinary perspective* (pp. 179-203). Cambridge, UK: Cambridge University Press.

19
20
21
22 Lewis, C., & Reading-Smith, M. (2008). Iraq: The war card. *The Center for Public*
23
24 *Integrity*. Retrieved December 23, 2009, from <http://projects.publicintegrity.org/WarCard/>

25
26
27 Magliano, J. P., & Radvansky, G. A. (2001). Goal coordination in narrative
28
29 comprehension. *Psychonomic Bulletin and Review*, 8, 372-376.

30
31
32 Nyhan, B., & Reifler, J. (2010). When corrections fail: The persistence of political
33
34 misperceptions. *Political Behavior*, 32, 303-330.

35
36
37 Oberauer, K., & Lewandowsky, S. (2008). Forgetting in immediate serial recall:
38
39 Decay, temporal distinctiveness, or interference? *Psychological Review*, 115, 544-576.

40
41
42 Payne, B. K., Lambert, A. J., & Jacoby, L. L. (2002). Best laid plans: Effects of goals
43
44 on accessibility bias and cognitive control in race-based misperceptions of weapons. *Journal*
45
46 *of Experimental Social Psychology*, 38, 384-396.

47
48
49 Pickel, K. L. (1995). Inducing jurors to disregard inadmissible evidence: A legal
50
51 explanation does not help. *Law and Human Behavior*, 19, 407-424.

52
53
54 Radvansky, G. A., & Copeland, D. E. (2001). Working memory and situation model
55
56 updating. *Memory & Cognition*, 29, 1073-1080.

57
58
59 Schul, Y., & Mazursky, D. (1990). Conditions facilitating successful discounting in
60
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.

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 Associates.

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.

Author Note

Preparation of this paper was facilitated by a Discovery Grant and an Australian Professorial Fellowship from the Australian Research Council to Stephan Lewandowsky.

Address correspondence to Ullrich Ecker or Stephan Lewandowsky at the School of Psychology (M304), University of Western Australia, Crawley, W.A. 6009, Australia.

Electronic mail may be sent to ullrich.ecker@uwa.edu.au or lewan@psy.uwa.edu.au. The lab's web site is located at <http://www.cogsciwa.com>. We thank Charles Hanich for research assistance.

For Review Only

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Footnotes

¹ The *3-MI/O-R* condition of Experiment 1, as well as the *L-MI/O-R* condition 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.

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

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

For Review Only

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Tables

Table 1

Contrasts Calculated on Inference Scores in Experiment 1

	Contrast	$t(154)$	p	d
0	<i>0-MI/0-R vs. 1-MI/3-R</i>	3.06	< .01	.49
1	<i>1-MI/0-R vs. 1-MI/1-R</i>	5.50	< .001	.89
2	<i>1-MI/0-R vs. 1-MI/3-R</i>	5.41	< .001	.87
3	<i>1-MI/1-R vs. 1-MI/3-R</i>	< 1		
4	<i>1-MI/0-R vs. 3-MI/0-R</i>	2.48	.01	.40
5	<i>3-MI/0-R vs. 3-MI/1-R</i>	5.09	< .001	.82
6	<i>3-MI/0-R vs. 3-MI/3-R</i>	8.25	< .001	1.33
7	<i>3-MI/1-R vs. 3-MI/3-R</i>	3.15	< .01	.51
8	<i>1-MI/1-R vs. 3-MI/1-R</i>	2.88	< .01	.46
9	<i>1-MI/3-R vs. 3-MI/3-R</i>	< 1		

Note. d , Cohen's d .

Table 2

Contrasts Calculated on Inference Scores in Experiment 2

	Contrast	<i>t</i> (132)	<i>p</i>	<i>d</i>
1	<i>L-MI/0-R vs. L-MI/L-R</i>	1.95	.05	.34
2	<i>L-MI/0-R vs. L-MI/noL-R</i>	1.65	.10	.29
3	<i>L-MI/L-R vs. L-MI/noL-R</i>	< 1		
4	<i>L-MI/0-R vs. noL-MI/0-R</i>	1.15	.25	.20
5	<i>noL-MI/0-R vs. noL-MI/L-R</i>	< 1		
6	<i>noL-MI/0-R vs. noL-MI/noL-R</i>	2.69	< .01	.47
7	<i>noL-MI/L-R vs. noL-MI/noL-R</i>	2.05	.04	.36
8	<i>L-MI/L-R vs. noL-MI/L-R</i>	2.45	.02	.43
9	<i>L-MI/noL-R vs. noL-MI/noL-R</i>	< 1		

Note. d, Cohen's d.

Figure Captions

Figure 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.


Note. *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.

Figure 2. Mean number of references to misinformation for all conditions in Experiment 2; error bars show standard errors of the mean.

Note. *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.


Figures

Figure 1.




1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Figure 2.




1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60



301x226mm (96 x 96 DPI)

View Only



291x218mm (96 x 96 DPI)

View Only

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60