



# Failure to accept retractions: A contribution to the continued influence effect

Andrea E. O’Rear<sup>1,2</sup> · Gabriel A. Radvansky<sup>1</sup>

Published online: 17 July 2019  
© The Psychonomic Society, Inc. 2019

## Abstract

Previous research has shown that when information about a narrative event is retracted, people continue to use that information even though it has been explicitly identified as incorrect. Not only can this occur for implicitly inferred information, but also when the change is stated explicitly. The current study explored whether this effect reflects, at least in part, an unwillingness of some readers to accept changes to their understanding. Experiment 1 assessed this using a continued influence effect paradigm with an additional probe asking whether participants believed the explicitly stated change. Most did not. Those that did accept it showed evidence of a reduced use of the incorrect information, while those that did not accept it performed similarly to those who received no correction (control). Experiment 2 included an additional explicit instruction that participants could say “don’t know” if they were unsure of how to respond. The pattern of results was largely the same as for Experiment 1. Experiment 3 modified the alternative account to increase plausibility, and added two additional stories/question sets to ensure effects were not limited to one set of materials. A greater number of participants found the retractions believable than in Experiments 1 and 2. Nonetheless, a similar pattern of results was found. Overall, these findings suggest that at least some of the evidence for the continued use of retracted information may be due to some people not accepting the retraction, even in the absence of external motivation to disregard it.

**Keywords** Retraction · Event model · Acceptance · Correction

## Introduction

People often make inferences while creating situation models of described events. Importantly, as a description unfolds, they may need to correct their understanding. This may be because they made inferences that turn out to be incorrect (e.g., Hamm & Hasher, 1992), or it may be explicitly noted that a previously given fact is incorrect and should be disregarded. This to-be-forgotten knowledge can still influence later responses. The aim of the current study is to explore to what degree an apparent inability to fully update one’s understanding is due to people not accepting a retraction as true.

## Change can be difficult

In early work by Kay (1955; see also Howe, 1970), people heard two passages and made repeated recalls separated by a week or more. After each attempt, people heard the passages again as an opportunity to correct any inaccurately recalled information. Strikingly, if people made recall errors early on, these errors persisted even after six recall-correction cycles. That is, their recall became very rigid, including the use of self-generated erroneous information, despite repeated presentations of the correct information. This finding has been found using cued recall (Amlund, Kardash, & Kulhavy, 1986), forced choice recognition (Howe, 1972), varied ordering of the repeated presentations and testing (Fritz, Morris, Bjork, Gelman, & Wickens, 2000), expository texts (Amlund et al., 1986; Fritz et al., 2000), attempts to direct attention to errors (Fritz et al., 2000), and for both reading and listening (Amlund et al., 1986). In all cases, the errors arose from inferences made during comprehension, but not when people were forced to guess or add information (Kang et al., 2011).

Of interest here is when people are told one thing that is later explicitly corrected, as with retractions (e.g., Ecker,

---

✉ Andrea E. O’Rear  
aorear@saintmarys.edu

<sup>1</sup> Department of Psychology, University of Notre Dame, Notre Dame, IN 46556, USA

<sup>2</sup> Department of Psychology, Saint Marys College, Notre Dame, IN 46556, USA

Lewandowsky, Swire, & Chang, 2011; Ecker, Lewandowsky, & Tang, 2010; Johnson & Seifert, 1994; Wilkes & Leatherbarrow, 1988; Wilkes & Reynolds, 1999). In the general paradigm, people read a series of messages about a scenario, such as a fire in a wholesaler building (Wilkes & Leatherbarrow, 1988). One message states a critical piece of information (e.g., police received reports that there were paint and gas cylinders in the room where the fire started). Then, a later message retracts it (e.g., police report that the room had actually been empty at the time of the fire). Finally, people are asked questions to determine whether they will use the original, retracted information (e.g., “Is there any evidence of careless management?”). The finding that people continue to make such inferences is called the Continued Influence Effect, or CIE.

The CIE is of concern because accurate understanding depends on the ability to modify knowledge when erroneous information has been identified. For example, Lewandowsky, Stritzke, Oberauer, and Morales (2005) reported that, years after the Iraq War, many Americans continued to claim that the invasion stopped Saddam Hussein from using weapons of mass destruction, despite news agencies repeatedly retracting this. Similarly, Greitemeyer (2014) reported that scientists continue to express belief in findings even when the reports of those findings have been retracted.

Several factors do not play a major role in the observation of the CIE. These include the use of emotional materials (Ecker, Lewandowsky, & Apai, 2011), natural aging (Guillory & Geraci, 2010), whether a retraction is presented immediately after a target statement or after a delay within the narrative (Johnson & Seifert, 1994), and whether people are socially prejudiced against the retraction (Ecker, Lewandowsky, Fenton, & Martin, 2014). The effect is not observed if the target information was mentioned in passing. It needs to be involved in the described event (Johnson & Seifert, 1994) to be part of the resulting situation model; otherwise, people are more likely to update their understanding (Wilkes & Leatherbarrow, 1988). The effect also persists, largely unchanged, if analyses are restricted to people who report remembering the retraction (Ecker, Lewandowsky, Swire, et al., 2011; Guillory & Geraci, 2010; Lewandowsky et al., 2005).

That said, there are factors that can change people’s understanding. First, if they are warned prior to reading that they may need to update their understanding, then they are more likely to do so (Ecker et al., 2010). Second, retractions are more effective if people are given an alternative explanation (Ecker et al., 2010; Ecker, Lewandowsky, & Apai, 2011; Johnson & Seifert, 1994). Third, if retrieval is delayed (e.g., 40 min) then people are more likely to use the more recent information (Ecker, Lewandowsky, & Apai, 2011). Fourth, people are more likely to update their understanding if the newer, more correct, information is repeatedly presented (Ecker, Lewandowsky, Swire, et al., 2011). Also, Wilkes and Leatherbarrow (1988) showed a smaller CIE if people are reminded of what they were told previously.

What has not been studied is if the size of the CIE may be due to people failing to accept the retraction. In such cases, the CIE would not be due to cognitive mechanisms operating in memory and comprehension, but to an unwillingness to use the retraction. The closest that this idea has been addressed is in a finding that people are more likely to change their understanding if the correct information comes from a more trustworthy source (Guillory & Geraci, 2013; but see Prasad et al., 2009). However, in the vast majority of studies there is no reason to distrust the source of the retraction. If anything, the opposite is true.

## Narrative updating

Although our focus is with updating prior incorrect information, there is a substantial literature on comprehension updating. For example, changes along an event dimension, such as a shift in location, time, character, goals, or a causal break (Zwaan & Radvansky, 1998), can result in increased reading times (Zwaan, Magliano, & Graesser, 1995) and decreased availability of prior event information in memory (e.g., Radvansky & Copeland, 2006). Still, the information about a prior event remains correct, although it refers to an earlier narrative time period, which differs from our concern here.

Another situation involves changes in more stable characteristics. For example, a story character can have a characteristic (e.g., vegetarian) changed to a new one (e.g., omnivore) (Kendeou, Smith, & O’Brien, 2013), an object changed from being unavailable to being available (O’Brien, Cook, & Guéraud, 2010), or an object changing states, such as a tree being cut down. In this research, comprehension is typically assessed in terms of reading times, with a slow-down occurring for sentences that are inconsistent with the originally described state or characteristic. This suggests that the outdated information is still activated to some degree. However, the prior information is still correct for that earlier narrative time period.

More in line with our concern about retracted information, Rapp and Kendeou (2007, 2009, see also Kendeou et al., 2013) had people read stories in which a character was described as acting consistent with a characteristic (e.g., being clumsy). Then, stories continued in one of three ways: (1) consistent (e.g., having trouble balancing); (2) inconsistent (e.g., always had great balance); or (3) inconsistent with an explanation (e.g., always had great balance but the floor was just waxed). Finally, people were shown an outcome sentence and asked if it was plausible. Reading times revealed activation of prior knowledge, but such an influence does not mean that people ultimately use the older information in their assessments. For example, ~90% of people normed in O’Brien et al. (2010) correctly answered questions about the availability/state of the critical objects (e.g., Was the tree cut

down?). Further, people in Rapp and Kendeou (2007, 2009) found outcomes most plausible when they followed whatever condition they were given. Thus, if the character was re-described as having great balance, people stuck with that rather than an original description of clumsiness. Outdated information is largely not used to make judgments, unlike the CIE. The type of updating is very different. In knowledge revision studies, the events described in the stories either changed over time (as with the tree being cut down; O'Brien et al., 2010) or led people to create inferences that turned out to be incorrect (as with a person only appearing clumsy; Rapp & Kendeou, 2007, 2009), while in the CIE studies, information was explicitly identified as incorrect at all times in the narrative world. The updating needed for retraction involves changing people's understanding of an event throughout narrative time – if they believe it. Thus, in both situations previous information can compete with newer information, but its status following updating is very different.

### Selective directed forgetting

The current study addresses the fate of knowledge that has been explicitly retracted in text, and thus should be disregarded. The long-term usage of discredited information is analogous to directed forgetting, in which people are told to disregard some information. The most common types of directed forgetting are item-based and list-based, in which people are told to forget information either immediately after reading it or after reading a set of items, respectively. However, neither of these captures the approach of CIE studies.

Instead, CIE studies have more in common with *selective directed forgetting* (Sahakyan, 2004). In these studies, people are presented with sets of information (e.g., sets of facts about Steve and Tom). Then, they are to selectively forget some information based on its content (e.g., forget everything about Steve). Like selective directed forgetting, people in CIE studies need to disregard information that is embedded in the larger set based on content. Selective directed forgetting is notoriously fickle, being observed under some circumstances (Delaney, Nghiem, & Waldum, 2009), but not in others (Akan & Sahakyan, 2018; Storm, Koppel, & Wilson, 2013). This may be mediated by the degree to which the targeted information is integrated with the rest (Pettijohn, 2016). Thus, our current study may have broader implications for the management of knowledge in long-term memory.

### Current study

The aim of the current study was to explore whether readers' failures to update understanding may be due to a lack of retraction acceptance. If so, mechanistic accounts of the CIE, such as automatic versus controlled retrieval processes (e.g.,

Ayers & Reder, 1998; Ecker, Lewandowsky, Swire, et al., 2011, Lewandowsky, Ecker, Seifert, Schwarz, & Cook, 2012) or event model updating (e.g., Ecker, Lewandowsky, Swire, et al., 2011; Johnson & Seifert, 1994; Lewandowsky et al., 2012), would also need to incorporate the role of belief.

Many studies of changes in understanding implicitly assume that participants take the information in the text at face value (apart from cases when instructed not to trust it; Ecker et al., 2010). To be consistent with prior studies, we used materials from Ecker et al. (2010) about a minibus accident with elderly passengers. In the retraction condition, police reported that it was not actually older adults on the minibus. In the alternative condition, additional information was provided that the passengers were actually hockey players. Finally, in the control condition there was no retraction. When older adults are mentioned in the retraction and alternative group responses to inference questions, relative to the control, this is the CIE. The CIE was larger for the retraction group than for the alternative group. Although Ecker et al. also explored the impact of warnings on the CIE, we did not do so here.

Here we explore the possibility that retractions may not be accepted, at least by some readers. They would continue to use the original information to make inferences about the described situation. Importantly, after asking questions to assess what sort of inferences people would draw, we also directly asked them if they believed the retraction.

## Experiment 1

The aim of Experiment 1 was to assess the rate of failure to accept a retraction, and whether the CIE would be reduced after removing people who did not accept it. However, if retracted information has an influence beyond its acceptance, based on cognitive mechanisms operating in comprehension and memory, then the CIE would remain unchanged. This was done using the paradigm and materials from Ecker et al. (2010) to ensure that our work lined up with what has already been done.

### Method

**Participants** Seventy-two people (48 female, age 18–22 years,  $M = 19.51$ ,  $SD = 1.08$ ) were recruited from the Department of Psychology participant pool at the University of Notre Dame. They were given partial course credit for their participation. All procedures were in compliance with the university's Institutional Review Board.

**Materials** The materials included the description of a minibus accident from Ecker et al. (2010), including the same inference, fact, and manipulation check questions. We added a

question at the end asking participants whether they believed the retraction (i.e., that it was not actually elderly people on the minibus). The sequence of messages, and the inference, fact, and manipulation check questions, are presented in Appendix 1.

**Procedure** First, participants were given a brief description of the study and then signed a consent form. They were then randomly assigned to one of three conditions: control, retraction, or alternative (i.e., retraction with an alternative account). People were tested separately on a computer where a series of messages was presented, one at a time, at people's own pace. After the last message, they were given a 6-min distractor task of answering questions about an autobiographical memory experience for a different study. At the end of the distractor period, participants were given the test questions. The first nine were the inference questions, which are labeled as such because participants must make an inference about the minibus passengers to answer the questions. If they answer these questions by making inferences conveying the idea that the passengers were elderly, it is an indication of the CIE. The next nine were the fact questions, to assess how well participants remember details of the story. The next two questions assessed memory for the retraction statement itself. Finally, the last question asked whether participants believed the retraction. Participants responded by typing an answer into a textbox below the question. Responses were scored as using the retracted target information when it was unambiguous and clear that they were using it to make an inference, and not just noting that it was mentioned. Examples of various types of responses made by

participants are presented in Table 1. Individual responses were independently scored by two researchers, Cohen's  $\kappa = .91$ . All discrepancies were resolved by discussion.

## Results and discussion

The data are reported in Table 2. For the inference questions, the primary dependent measure is the rate at which the idea that the passengers were older adults was used, even though this was retracted. The inference data are shown in Fig. 1. There was a main effect of Condition,  $F(2,69) = 6.20$ ,  $p = .003$ ,  $\eta_p^2 = .15$ . Critical inferences were made more often in the control condition than in the retraction and alternative conditions,  $F(1,46) = 4.40$ ,  $p = .04$ ,  $\eta_p^2 = .09$  and  $F(1,46) = 12.19$ ,  $p = .001$ ,  $\eta_p^2 = .21$ , respectively, which did not significantly differ,  $F(1,46) = 2.12$ ,  $p = .15$ ,  $\eta_p^2 = .04$ . The probability of giving a report that included target inferences was greater than zero for the control group,  $t(23) = 10.69$ ,  $p < .001$ ,  $d = 2.18$ ; the retraction group,  $t(23) = 6.91$ ,  $p < .001$ ,  $d = 1.41$ ; and the alternative group,  $t(23) = 4.26$ ,  $p < .001$ ,  $d = .87$ . Overall, while retraction reduced the rate of using the target information, providing an alternative account did not further reduce this influence.

Next, we removed participants in the retraction and alternative conditions who forgot the retraction, because if they forgot it, it could not be used. There was still a main effect of Condition,  $F(2,59) = 3.79$ ,  $p = .03$ ,  $\eta_p^2 = .11$ . People made critical inferences marginally more often in the control than in the retraction condition,  $F(1,42) = 3.25$ ,  $p = .08$ ,  $\eta_p^2 = .07$ , and significantly more often than in the alternative condition,  $F(1,40) = 7.34$ ,  $p = .01$ ,  $\eta_p^2 = .15$ . The last two did not differ,  $F(1,36) = 0.86$ ,  $p = .36$ ,  $\eta_p^2 = .02$ . Again, the probability of

**Table 1** Example responses for each type of response classification and type of question asked

Question asked	Response type	Example responses
4. Why do you think it was difficult getting the uninjured passengers up the embankment?	Critical inference	Considering the old age of the passengers, I would assume that they would struggle to climb the embankment
	No critical inference	The embankment was steep and even uninjured passengers would be in a state of shock
	Inference included but not counted	Climbing up an embankment is hard if the people were tired from playing a long hockey game. If they were elderly, it is also probably hard because again, they are just old
10. On which day did the accident occur?	Correct fact	Sunday
	Incorrect fact	Bingo day?
20. Were you aware of any corrections in the messages that you read?	Aware	There appeared to be confusion in the report about whether the passengers were elderly or not
	Unaware	No
21. Do you believe any corrections in the messages that you read, or did you think that they were wrong?	Believed	I believed the correction
	Did not believe	I think that they were wrong, because some of the other statements make it [plausible] that the passengers were elderly (trouble telling injured from uninjured, trouble evacuating them, trouble getting them up the embankment, etc.)

**Table 2** Responses to the questions asked in Experiment 1

	Condition	N	Inference	Fact	Memory	Belief
All data	Control	24	.48 (.05)	.75 (.03)		
	Retraction	24	.34 (.05)	.77 (.03)	.83 (.08)	.33 (.10)
	Alternative	24	.23 (.05)	.68 (.04)	.75 (.09)	.42 (.10)
Retraction remembered	Retraction	20	.35 (.06)	.78 (.04)	1.00 (.00)	.25 (.10)
	Alternative	18	.27 (.06)	.71 (.04)	1.00 (.00)	.39 (.12)
Retraction believed	Retraction	5	.20 (.08)	.82 (.06)	1.00 (.00)	1.00 (.00)
	Alternative	7	.03 (.07)	.71 (.03)	1.00 (.00)	1.00 (.00)
Retraction not believed	Retraction	15	.40 (.07)	.77 (.05)	1.00 (.00)	0.00 (.00)
	Alternative	11	.42 (.07)	.71 (.07)	1.00 (.00)	0.00 (.00)

For Inference and Fact questions, these are the proportion correct, for the Memory questions, this is whether a person indicated memory for the retraction on either or both of the questions, and for the Belief question, this is whether or not a person believed the retraction (if remembered). These last two are not defined for the control condition

Standard errors are in parentheses

using the target information was greater than zero for the retraction group,  $t(19) = 6.31, p < .001, d = 1.41$ ; and the alternative group,  $t(17) = 4.22, p < .001, d = 1.00$ . Thus, as in previous research (Ecker, Lewandowsky, Swire et al., 2011; Guillory & Geraci, 2010; Lewandowsky et al., 2005), there was no evidence that the CIE was much influenced by forgetting the retraction.

Next, we explored whether the magnitude of the CIE is influenced by people who do not accept the retraction. For this analysis we further removed participants in the retraction and alternative conditions who indicated they did not believe the retraction. Of critical importance, as shown in Table 1, this was the majority of participants who remembered the retraction. There was a main effect of Condition,  $F(2,33) = 15.64, p < .001, \eta_p^2 = .49$ . Inferences using the target information were made more often in the control than in the retraction and alternative conditions,  $F(1,27) = 6.95, p = .01, \eta_p^2 = .20$  and  $F(1,29) = 27.30, p < .001, \eta_p^2 = .48$ , respectively, which marginally differed from one another,  $F(1,10) = 4.70, p = .06, \eta_p^2 = .32$ . Thus, providing participants with an alternative can further reduce the CIE. For these data, the probability of using target inferences was trending towards being greater than zero for the retraction group,  $t(4) = 2.45, p = .07, d = 1.10$ , and it was not significant for the alternative group,  $t(6) = 1.00, p = .36, d = .38$ . Thus, the CIE was reduced when only participants who accepted the retraction were analyzed.<sup>1</sup> This suggests that retractions may be more effective at isolating the influence of the retracted material from performance if people accept the retraction. That is, the CIE may be inflated by

<sup>1</sup> It may be argued that the sample size is too small to conclude whether there is a difference from zero. That may be true, but it is also beside the point. The point we are making is that there is a large number of people who do not accept the retraction (i.e., the disbelief rate is high), and this is having a major influence on the CIE magnitude.


including participants who do not accept the retraction and continue using that information to make inferences.

As a final step, to show that when the retraction is not accepted, it is as if there was no retraction, we compared the inference data for the control condition with the retraction and alternative conditions for participants who remembered, but did not accept, the retraction. As shown in Fig. 2, there was no effect of condition,  $F(2,47) = 0.55, p = .58, \eta_p^2 = .02$ . Thus, when a retraction was not accepted, the misinformation continued to be used.

For the fact questions (see Fig. 3), the main effect of Condition was not significant,  $F(2,69) = 2.13, p = .13, \eta_p^2 = .06$ . All three groups were similarly able to answer questions about explicitly provided information. This was also true with the removal of people who forgot the retraction,  $F(2,59) = 1.02, p = .37, \eta_p^2 = .03$ , and those who did not accept it,  $F(2,33) = 1.04, p = .37, \eta_p^2 = .06$ . Thus, these factors had no major influence on memory for the information read in the messages.

## Experiment 2

The aim of Experiment 2 was to replicate the findings of Experiment 1 and to explore the possible influence of a latent demand characteristic. Specifically, by asking people to provide answers to questions, there may be an implicit demand to provide some answer, even if it is known that the information was retracted. Work in eyewitness identification has found an implicit demand to choose a person from a line-up, resulting in false identifications. However, when people are explicitly told that the perpetrator may not be in the line-up and that they can give a “not there” response, false identifications diminish (e.g., Malpass & Devine, 1981; Steblay,



**Fig. 1** The proportion of critical inferences made in Experiment 1 by each group, further divided into those who stated they remembered the retraction and those who stated they believed the retraction. Error bars represent standard error of the mean

1997). Using this as inspiration, we sought to lessen the demand to report something (and use retracted information known to be incorrect), so we explicitly told people that they could respond “don’t know” when they did not feel that they could provide an answer.


## Method

**Participants** Seventy-two people (41 female, age 18–36 years,  $M = 20.37$ ,  $SD = 2.54$ ) were recruited from the Department of Psychology participant pool at the University of Notre Dame. They were given partial course credit for their participation. All procedures were in compliance with the university’s Institutional Review Board.


**Materials and procedure** The materials and procedure for Experiment 2 were the same as those in Experiment 1, with the addition of a sentence at test explicitly stating that a “don’t know” response could be used. Again, responses were independently scored by two researchers, Cohen’s  $\kappa = .91$ , and all discrepancies were resolved by discussion.

## Results and discussion

The number of participants who gave a “don’t know” or “don’t remember” response at least once in Experiment 2 was 54, as compared to 31 in Experiment 1. Consistent with other studies of a bias to avoid “don’t know” responses when participants are not explicitly told that they may do so, instances of such responses occurred less often in Experiment 1 (not explicitly told;



**Fig. 2** The proportion of critical inferences made in Experiment 1 by only those who did not believe the retraction, compared to control. Error bars represent standard error of the mean



**Fig. 3** The proportion of correct responses to fact questions made in Experiment 1 by each group, further divided into those who stated they remembered the retraction and those who stated they believed the retraction. Error bars represent standard error of the mean

$M = 0.90$ ) than in Experiment 2 (explicitly told;  $M = 1.92$ ),  $F(1,142) = 13.08$ ,  $p < .001$ ,  $\eta_p^2 = .08$ . To determine if this influenced the CIE, we turn to the testing data.

The data are reported in Table 3. For inference questions (data shown in Fig. 4), there was a main effect of Condition,  $F(2,69) = 16.12$ ,  $p < .001$ ,  $\eta_p^2 = .32$ . The target information was used in inferences more often in the control than in the retraction and alternative conditions,  $F(1,46) = 12.81$ ,  $p = .001$ ,  $\eta_p^2 = .22$  and  $F(1,46) = 34.03$ ,  $p < .001$ ,  $\eta_p^2 = .43$ , respectively, which were marginally different from one another,  $F(1,46) = 3.11$ ,  $p = .09$ ,  $\eta_p^2 = .06$ . The probability of giving a report that included critical inferences was greater than zero for the control group,  $t(23) = 9.84$ ,  $p < .001$ ,  $d = 2.01$ ; the retraction group,  $t(23) = 4.46$ ,  $p < .001$ ,  $d = .91$ ; and the alternative group,  $t(23) = 2.87$ ,  $p = .009$ ,  $d = .59$ . Again, while retraction reduced the rate of using the target information, providing an alternative explanation did not substantially reduce this. Moreover, in all cases, the retracted information was used to some degree.

Next, we removed participants in the retraction and alternative conditions who forgot the retraction. There was a main effect of Condition,  $F(2,62) = 14.88$ ,  $p < .001$ ,  $\eta_p^2 = .32$ . Inferences using the target information were made more often in the control than in the retraction and alternative conditions,  $F(1,43) = 11.85$ ,  $p = .001$ ,  $\eta_p^2 = .22$  and  $F(1,42) = 30.82$ ,  $p < .001$ ,  $\eta_p^2 = .42$ , respectively, which did not differ,  $F(1,39) = 2.78$ ,  $p = .10$ ,  $\eta_p^2 = .07$ . Again, the probability of including target information was greater than zero in the retraction group,  $t(20) = 4.06$ ,  $p < .001$ ,  $d = .89$ ; and the alternative group,  $t(19) = 2.36$ ,  $p = .029$ ,  $d = .53$ . As in Experiment 1 and previous research, memory for the retracted information had little influence on the CIE.

Next, we dropped any participants in the retraction and alternative conditions who did not accept the retraction. As shown in Table 3, this was nearly half of the people who remembered the retraction in these two conditions, further reinforcing the idea that many people did not accept the

retraction. For the inference data, there was a main effect of Condition,  $F(2,43) = 28.49$ ,  $p < .001$ ,  $\eta_p^2 = .57$ . The target information was used more often in the control than in the retraction and alternative conditions,  $F(1,33) = 38.16$ ,  $p < .001$ ,  $\eta_p^2 = .54$  and  $F(1,33) = 27.27$ ,  $p < .001$ ,  $\eta_p^2 = .45$ , respectively, which did not differ,  $F(1,20) = .24$ ,  $p = .63$ ,  $\eta_p^2 = .01$ . For these data, the probability of giving a report that used the target information was only trending greater than zero for the retraction group,  $t(10) = 1.94$ ,  $p = .08$ ,  $d = 0.59$ , and was not greater than zero for the alternative group,  $t(10) = 1.00$ ,  $p = .34$ ,  $d = 0.30$ . Thus, it appears that having the option to say “don’t know” reduced the tendency to use the retracted information. The rate of using the retracted information in Experiment 2 was compared to the rate in Experiment 1, and it was indeed found to be reduced,  $F(1,142) = 4.47$ ,  $p = .04$ ,  $\eta_p^2 = .03$ .

Finally, to show that not accepting the retraction is equivalent to never getting it, we compared the data for participants not accepting the retraction in the retraction and alternative conditions with the control condition. This analysis, shown in Fig. 5, revealed an effect of Condition,  $F(2,40) = 6.38$ ,  $p = .004$ ,  $\eta_p^2 = .24$ , which is due to a lower use of the retracted information in the alternative condition (control vs. alternative:  $F(1,31) = 11.73$ ,  $p = .002$ ,  $\eta_p^2 = .27$ ; retraction vs. alternative:  $F(1,17) = 8.31$ ,  $p = .01$ ,  $\eta_p^2 = .33$ ; control vs. retraction:  $F(1,32) < 1$ ,  $\eta_p^2 = .005$ ). Thus, when people do not accept the retraction, it may have a reduced influence, but only if an alternative is given.<sup>2</sup>

<sup>2</sup> When considering why people in the alternative condition in Experiment 2 who did not believe the change would make fewer inferences based on retracted information than those in Experiment 1, we thought that the higher rate of “don’t know” responses in Experiment 2 could lead to fewer inferences being made. We compared the two alternative/don’t believe conditions and found that those in Experiment 2 ( $M = 8.22$ ) made fewer inferences than those in Experiment 1 ( $M = 8.91$ ),  $F(1,18) = 4.97$ ,  $p = .04$ ,  $\eta_p^2 = .22$ .

**Table 3** Responses to the questions asked in Experiment 2

	Condition	N	Inference	Fact	Memory	Belief
All data	Control	24	.53 (.05)	.76 (.04)		
	Retraction	24	.25 (.06)	.63 (.03)	.88 (.07)	.54 (.10)
	Alternative	24	.13 (.04)	.73 (.03)	.83 (.08)	.54 (.10)
Retraction remembered	Retraction	21	.25 (.06)	.63 (.03)	1.00 (.00)	.52 (.11)
	Alternative	20	.12 (.05)	.73 (.03)	1.00 (.00)	.55 (.11)
Retraction believed	Retraction	11	.03 (.02)	.62 (.04)	1.00 (.00)	1.00 (.00)
	Alternative	11	.06 (.06)	.70 (.04)	1.00 (.00)	1.00 (.00)
Retraction not believed	Retraction	10	.49 (.07)	.65 (.05)	1.00 (.00)	0.00 (.00)
	Alternative	9	.19 (.08)	.77 (.05)	1.00 (.00)	0.00 (.00)

For Inference and Fact questions, these are the proportion correct; for the Memory questions, this is whether a person indicated memory for the retraction on either or both of the questions; and for the Belief question, this is whether or not a person believed the retraction (if remembered). These last two are not defined for the control condition

Standard errors are in parentheses


For the fact questions (data shown in Fig. 6), there was a main effect of Condition,  $F(2,69) = 4.07$ ,  $p = .02$ ,  $\eta_p^2 = .11$ . This was also true with the removal of participants who forgot the retraction,  $F(2,62) = 3.66$ ,  $p = .03$ ,  $\eta_p^2 = .11$ . This is largely due to lower accuracy in the retraction group. This small effect may be due to some people having trouble managing the retraction along with other information. This does not appear in other research, and it is unclear how it would bear on our theoretical understanding of retraction. That said, it should be noted that this effect was only marginally significant with the removal of those who did not accept the retraction,  $F(2,43) = 2.68$ ,  $p = .08$ ,  $\eta_p^2 = .11$ , which may be attributable to the lower power of this sample.

### Experiment 3

The aim of Experiment 3 was to replicate Experiments 1 and 2 and to address several factors. One concern about


Experiments 1 and 2 is that it may not seem plausible that rescue workers would confuse older adults with hockey players, hence the low levels of belief in the alternative condition. For Experiment 3, we revised this description so the alternative was multiple sclerosis patients. We also added two new scenarios. One was a modification of the fire scenario of Wilkes and Leatherbarrow (1988), and the other was a new description we created about a spinach recall. For these, we strove to create descriptions in which the two alternatives were plausible (i.e., gas cans vs. fireworks for the fire scenario, and expired pesticides vs. dirty factory equipment for the tainted spinach scenario).

Another factor we explored was how the information was presented. Specifically, in Experiments 1 and 2, like most retraction studies, the descriptions were presented as a series of messages, similar to a list-learning paradigm. However, this is unlike the way that most events in the world are communicated. A news report is not a series of messages – it is a



**Fig. 4** The proportion of critical inferences made in Experiment 2 by each group, further divided into those who stated they remembered the retraction and those who stated they believed the retraction. Error bars represent standard error of the mean





**Fig. 5** The proportion of critical inferences made in Experiment 2 by only those who did not believe the retraction, compared to control. Error bars represent standard error of the mean

narrative. People process information differently depending on whether they view it as a narrative or not (Zwaan, 1994). In real-world cases, when people hear information that might later be retracted, it is likely to be encountered as a narrative. To address the impact, if any, of presentation format, half of the participants in Experiment 3 received the scenarios as lists of messages, and the other half as narratives.

Beyond this, we made two procedural changes. One was to use a repeated-measures design. In Experiments 1 and 2 each participant received the control, retraction, or alternative version of the minibus scenario. With the two new scenarios added in this experiment, each person would read one description for each condition. This provided an opportunity to replicate our findings using a different paradigm and to extend our findings to new materials.


We also manipulated what was the original and the alternative. In Experiments 1 and 2, the original information was

always that the passengers were older adults and the alternative was always that they were hockey players. In Experiment 3, which option was the original and which was the alternative was counterbalanced across participants for all three of the passages.

**Method**

**Participants** Sixty-four people (28 female, age 18–25 years  $M = 19.52, SD = 1.35$ ) were recruited from the Department of Psychology participant pool at the University of Notre Dame. They were given partial course credit for their participation. All procedures were in compliance with the university’s Institutional Review Board.

**Materials** The materials included three experimental and three filler scenarios. One experimental scenario was a modified



**Fig. 6** The proportion of correct responses to fact questions made in Experiment 2 by each group, further divided into those who stated they remembered the retraction and those who stated they believed the retraction. Error bars represent standard error of the mean

version of that used in Experiments 1 and 2, with the alternative changed from hockey players to multiple sclerosis patients. Another was adapted from Wilkes and Leatherbarrow (1988), and the alternative of fireworks and sparklers was added. The third was newly written about a spinach recall. This scenario is presented in Appendix 2. Each experimental scenario had control, retraction, and alternative versions. For all scenarios: (1) which option was the original and which was the alternative was counterbalanced across participants, and (2) there was a version that was written as a series of separate messages, as with most work in this area, and another that was written as a narrative. Finally, there were three filler scenarios to help obscure the retraction manipulation. These were similar in length to the experimental scenarios.

**Procedure** First, reading instructions were provided to participants. These instructions varied for the messages versus narratives conditions, as shown in Appendices 1 and 2, to emphasize the condition. Next, all six scenarios were read, with the first and last ones always being the same filler scenarios, and the remaining four presented in a different random order for each participant. At the beginning of each scenario, a title was provided. These scenarios were presented one message/sentence at a time. After all scenarios were read, people were given the inference, fact, and manipulation check questions for the three experimental stories. At the beginning of each question set, the title of the relevant scenario was presented as an indicator for the participant. Participants were explicitly told that they could respond “don’t know” if they did not know an answer. Individual responses were independently scored by two researchers, Cohen’s  $\kappa = .81$ , and any discrepancies were resolved by discussion.

## Results and discussion

The data are reported in Table 4. The data for each question type were submitted to a 2 (Material Type: List or Narrative)  $\times$  3 (Condition: Control, Retraction, or Alternative) mixed ANOVA. For the inference questions (data shown in Fig. 7), there was a main effect of Condition,  $F(2,124) = 24.08$ ,  $p < .001$ ,  $\eta_p^2 = .28$ . People made critical inferences more often in the control condition than in the retraction and alternative conditions,  $F(1,62) = 8.12$ ,  $p = .006$ ,  $\eta_p^2 = .12$  and  $F(1,62) = 42.30$ ,  $p < .001$ ,  $\eta_p^2 = .41$ , respectively, which also differed from one another,  $F(1,62) = 19.65$ ,  $p < .001$ ,  $\eta_p^2 = .24$ . The probability of giving a report that included inappropriate inferences was greater than zero for the control group,  $t(63) = 14.40$ ,  $p < .001$ ,  $d = 1.80$ ; the retraction group,  $t(63) = 10.31$ ,  $p < .001$ ,  $d = 1.29$ ; and the alternative group,  $t(63) = 6.16$ ,  $p < .001$ ,  $d = .77$ . Neither the main effect of Material Type nor the interaction were significant,  $F < 1$ ,  $F(2,124) = 1.54$ ,  $p = .22$ ,  $\eta_p^2 = .02$ , respectively. Thus, presentation format did not appear to have an impact (all figures shown are collapsed across

Material Type). The retraction of the target information reduced the rate at which it was used to answer the inference questions. Receiving an alternative explanation further reduced this influence, perhaps because these alternatives were more plausible. Finally, in all cases, people used the retracted information to some degree.

Following this analysis, we removed participants in the retraction and alternative conditions who forgot the retraction. Data were only removed for the affected conditions, not an entire participant. For example, if a person remembered the retraction for the minibuss story but not the fire story, their data would only be removed from the fire story condition. After this, people made inappropriate inferences more often in the control condition than in the retraction and alternative conditions,  $F(1,32) = 6.13$ ,  $p = .02$ ,  $\eta_p^2 = .16$  and  $F(1,39) = 17.92$ ,  $p < .001$ ,  $\eta_p^2 = .31$ , respectively, which did not significantly differ from one another,  $F(1,23) = 2.97$ ,  $p = .10$ ,  $\eta_p^2 = .11$ . Again, the probability of giving a report that included critical inferences was greater than zero for the retraction group,  $t(32) = 7.75$ ,  $p < .001$ ,  $d = 1.35$ ; and the alternative group,  $t(39) = 5.73$ ,  $p < .001$ ,  $d = .91$ . Thus, as in previous research, there was no evidence that the CIE was due to the inclusion of people who forgot the retraction.

Having again replicated the finding that the CIE is not due to poor retraction memory, we then explored whether it is affected by a failure to accept the retraction. For this analysis we removed people in the retraction and alternative conditions who said they did not believe the retraction. Again, data were only removed for the affected conditions. As can be seen in Table 4, this was a large amount of people, although not quite as large as was seen in the previous experiments. This may be due to the updated materials. After this selection, people made critical inferences more often in the control condition than in the retraction, but not significantly so,  $F(1,13) = 2.98$ ,  $p = .11$ ,  $\eta_p^2 = .23$ , possibly due to lower power after dropping so many people. However, people made these inferences more often in the control condition than in the alternative condition,  $F(1,26) = 13.94$ ,  $p = .001$ ,  $\eta_p^2 = .35$ . The retraction and alternative conditions did not differ from one another,  $F < 1$ . For these data, the probabilities of giving a report that included critical inferences were still greater than zero for the retraction group,  $t(14) = 4.25$ ,  $p < .001$ ,  $d = 1.10$ ; and the alternative group,  $t(27) = 4.22$ ,  $p < .001$ ,  $d = 0.80$ . As in Experiments 1 and 2, the CIE decreased when only those people who remembered and believed the retraction were analyzed. However, in this experiment, perhaps because of changes in the materials and the design, a CIE was observed even after people who had forgotten and/or did not believe the retraction were set aside.


As a final step, we assessed performance when people reported not believing the retraction compared to the control condition. When the inference data were assessed after this

**Table 4** Responses to the questions asked in Experiment 3


	Condition	N	Inference	Fact	Memory	Belief
All data	Control (list)	32	.48 (.04)	.59 (.03)		
	Control (narrative)	32	.38 (.05)	.57 (.04)		
	Control (mean)	64	.43 (.03)	.58 (.03)		
	Retraction (list)	32	.31 (.04)	.54 (.03)	.63 (.09)	.44 (.09)
	Retraction (narrative)	32	.32 (.05)	.54 (.03)	.41 (.09)	.41 (.09)
	Retraction (mean)	64	.31 (.06)	.54 (.02)	.52 (.06)	.42 (.06)
	Alternative (list)	32	.15 (.03)	.52 (.04)	.56 (.09)	.56 (.09)
	Alternative (narrative)	32	.17 (.04)	.56 (.04)	.69 (.08)	.59 (.09)
	Alternative (mean)	64	.16 (.03)	.54 (.03)	.63 (.06)	.58 (.06)
Retraction remembered	Retraction (list)	20	.33 (.05)	.59 (.03)	1.00 (.00)	.45 (.11)
	Retraction (narrative)	13	.32 (.08)	.59 (.05)	1.00 (.00)	.46 (.14)
	Retraction (mean)	33	.33 (.04)	.59 (.03)	1.00 (.00)	.46 (.09)
	Alternative (list)	18	.18 (.05)	.53 (.05)	1.00 (.00)	.72 (.11)
	Alternative (narrative)	22	.18 (.04)	.58 (.05)	1.00 (.00)	.68 (.10)
	Alternative (mean)	40	.18 (.03)	.56 (.03)	1.00 (.00)	.70 (.07)
Retraction believed	Retraction (list)	9	.30 (.09)	.59 (.05)	1.00 (.00)	1.00 (.00)
	Retraction (narrative)	6	.17 (.06)	.65 (.08)	1.00 (.00)	1.00 (.00)
	Retraction (mean)	15	.24 (.06)	.62 (.04)	1.00 (.00)	1.00 (.00)
	Alternative (list)	13	.12 (.05)	.53 (.05)	1.00 (.00)	1.00 (.00)
	Alternative (narrative)	15	.17 (.05)	.59 (.06)	1.00 (.00)	1.00 (.00)
	Alternative (mean)	28	.15 (.04)	.57 (.04)	1.00 (.00)	1.00 (.00)
Retraction not believed	Retraction (list)	11	.36 (.06)	.60 (.05)	1.00 (.00)	0.00 (.00)
	Retraction (narrative)	7	.44 (.12)	.54 (.07)	1.00 (.00)	0.00 (.00)
	Retraction (mean)	18	.40 (.06)	.57 (.04)	1.00 (.00)	0.00 (.00)
	Alternative (list)	5	.33 (.09)	.52 (.14)	1.00 (.00)	0.00 (.00)
	Alternative (narrative)	7	.19 (.08)	.56 (.06)	1.00 (.00)	0.00 (.00)
	Alternative (mean)	12	.25 (.06)	.55 (.06)	1.00 (.00)	0.00 (.00)

For Inference and Fact questions, these are the proportion correct; for the Memory questions, this is whether a person indicated memory for the retraction on either or both of the questions; and for the Belief question, this is whether or not a person believed the retraction (if remembered). These last two are not defined for the control condition

Standard errors are in parentheses



**Fig. 7** The proportion of critical inferences made in Experiment 3 by each group, further divided into those who stated they remembered the retraction and those who stated they believed the retraction. Error bars represent standard error of the mean




**Fig. 8** The proportion of critical inferences made in Experiment 3 by only those who did not believe the retraction, compared to control. Error bars represent standard error of the mean

selection, as shown in Fig. 8, people made critical inferences at similar rates in the control and retraction conditions,  $F(1,16) = 2.68, p = .12, \eta_p^2 = .14$ . There was a trend for people to make these inferences more often in the control than in the alternative condition, although the power was very low,  $F(1,10) = 4.57, p = .06, \eta_p^2 = .31$ . Overall, when people did not accept the retraction, it was still having an influence. More importantly, as with Experiments 1 and 2, it can be seen that there is a great deal of hesitancy by participants to accept the retractions when they were provided.

For the fact questions (data shown in Fig. 9), neither the main effect of Material Type, Condition, nor the interaction were significant, all  $F_s < 1$ . Thus, neither the nature of the materials (list or narrative) nor the presence or absence of retraction had any influence on general memory for the materials that were explicitly read.

## General discussion

The current study assessed whether part of the reason that people use retracted information is a failure to accept a retraction. Across our experiments, retraction acceptance was quite low. Without taking this into account, the magnitude of the CIE was larger. When data were assessed based on retraction acceptance, those who accepted it showed a reduced CIE. Those who did not accept it generally used the inappropriate information at a similar rate as people who never got a retraction. Thus, retraction acceptance is an important component in the magnitude of the CIE. Further, Experiment 2 found that explicitly telling people they can reply “don’t know” reduced the CIE rate. This is in line with research showing that this manipulation influences memory reports in other areas (e.g., Malpass & Devine, 1981; Steblay, 1997). Experiment 3



**Fig. 9** The proportion of correct responses to fact questions made in Experiment 3 by each group, further divided into those who stated they remembered the retraction and those who stated they believed the retraction. Error bars represent standard error of the mean

increased generalizability by showing that a large proportion of people (over 40% of those who remembered the retraction) did not believe the retraction across different sets of materials, although a larger CIE remained than in the other experiments. Possibilities for this difference are further discussed in the [Limitations](#) section.

Finally, Experiment 3 did not find any influence of whether information was presented as a list of messages or a narrative. Zwaan (1994) found that how people view what they are reading can influence the mental representations that are used and the inferences that are drawn. The absence of this pattern here may be because people were treating texts as descriptions of events, and not as expository information, regardless of presentation format. Although one could argue that presenting a narrative one sentence at a time could have led people to view the information as a list of messages, this is unlikely. Previous narrative comprehension work has presented narratives one sentence at a time, with success, often to allow for the assessment of reading rates (e.g., Radvansky & Copeland, 2010; Zwaan, Magliano, & Graesser, 1995). As such studies have found distinctive effects of narrative structure, it is reasonable to assume that presenting sentences individually did not prevent people from interpreting the narratives.

How do the current findings bear on prior findings? One likely explanation for the continued use of outdated information is based on event model updating (e.g., Ecker, Lewandowski, Swire, et al., 2011; Johnson & Seifert, 1994; Lewandowski et al., 2012). After people create event models of a described scenario(s), when a retraction is made, people need to update their event models. If no alternative explanation is provided, people are more likely to use the retracted information rather than leave a part missing from their model. That is, people are uncomfortable with incomplete event models, so using retracted information is preferable. If an alternative is provided, people are less likely to use the retracted information because they can re-form a more complete event model. However, if the alternative is not accepted, the CIE can still occur.

### Implications for research on the CIE

We turn next to roles played by specific manipulations. For the study by Ecker et al. (2010) that served as the basis for our study, one manipulation was a warning that some information might not be reliable. This warning was either general (i.e., sometimes facts are not checked before news reports are published) or specific (i.e., a definition of the CIE with examples). Regardless of warning type, the CIE persisted. For the general warning, performance was the same as with retraction only, but worse than when an alternative was provided. This warning may not be compelling enough for people to accept a retraction. It could also be interpreted that the original

information was correct and that the retraction was unreliable. It is up to participants to decide which information to accept.

For the specific warning, performance was better than with retraction alone, but no different from when an alternative was given. Giving both a specific warning and an alternative improved performance further. Having both the alternative and a specific warning gives people two alternatives, thereby increasing effectiveness. However, it is still possible that some people may not accept the alternative because the idea that the minibus passengers were older adults better fit their understanding than hockey players.

Finally, Ecker et al. (2010) found that with specific warnings, there was no difference in the number of references to the retracted information between participants who said they did versus did not believe the retraction. This is in contrast to what we have found here. One possible factor is that Ecker et al. assessed belief using only those people who used the retracted information at least once. In our Experiments 1 and 2 (those using only Ecker et al.'s materials), 16 and 27 participants, respectively, did not use the retracted information. Of those subsets, 11 and 22, respectively, believed the retraction. By excluding people who did not make any inappropriate inferences, the comparison of *did* versus *did not believe* may be skewed. We also did not include any warnings. As noted by Ecker et al., including warnings influences how participants encode the information and may affect whether a retraction is to be accepted or not.

Our study also has some affinity with work by Guillory and Geraci (2013), who showed that retraction effectiveness was influenced by the trustworthiness of the retraction source. When the source was more trustworthy, the retraction was more effective. Similarly, in the refutation literature, Sparks and Rapp (2011) found that people rely on source credibility when making judgments about the likelihood of characters' future behaviors that were either consistent or inconsistent with traits others had described those characters possessing. Although Guillory and Geraci (2013) asked about belief, the question was framed in terms of the credibility of the retraction source. Thus, while that study provides some evidence that belief can impact the CIE, its focus was on source trustworthiness rather than specifically assessing acceptance.

While other studies have examined the role of belief in misinformation persistence (e.g., Lewandowsky, Oberauer, & Gignac, 2013; Lewandowsky et al., 2005; Prasad et al., 2009), they tend to focus on pre-existing opinions and beliefs. For example, Prasad et al. (2009) interviewed supporters of President Bush who indicated that they believed in a link between Saddam Hussein and the 9/11 terrorist attack. The interviewers showed those people a quote from President Bush stating that there was no such link. Despite this, 80% of the people persisted in their belief of a connection. Although studies such as this are important for understanding the role of existing beliefs, we are hesitant to draw

comparisons with our own study, which did not assess this. What is interesting is the sheer low rate of retraction acceptance in the likely absence of such pre-experimental beliefs, or their direct manipulation, on the issue discussed in the scenario.

One study of the role of belief within an experimental setting was by Rich and Zaragoza (2016). They examined the CIE when the original misinformation was presented explicitly versus implicitly, finding a stronger CIE for implicit information. From a post-test questionnaire, they found that fewer than half of participants said they believed the correction, consistent with our findings. They then tested whether people in the implicit condition were less likely to believe the correction. No difference in belief rates was found, suggesting that there was some reason other than belief behind the greater CIE in the implicit condition.

Importantly, in Rich and Zaragoza (2016) the misinformation was that the police suspected the homeowners' son was involved in the theft, and the correction was that they later learned the son was out of town that day. This is more in line with knowledge revision and refutation studies (e.g., Rapp & Kendeou, 2007, 2009) in that both pieces of information were true at some point in the narrative world. The reader was not told that the son committed the crime, only that he was a suspect. Therefore, that assessment of belief is for a different type of updating than what was used in our experiments.

### Other lines of study

Taking a different perspective, the knowledge revision components framework (KReC) by Kendeou and O'Brien (2014; as cited in Kendeou, Butterfuss, van Boekel, & O'Brien, 2017) describes how old and new information can compete during comprehension leading to an observed influence of outdated information. They argue that previously learned information is not removed from memory and can be reactivated. To reduce interference from this reactivated information, people first need to integrate the new information with the old. Then to revise one's knowledge, more instances of the updated information need to be encoded to outcompete the older information for activation. In terms of the present study, it could be that one's lack of belief in a retraction may prevent them from properly integrating the incorrect and correct information because they do not find the correction to be relevant. As a result, the CIE persists.

One way knowledge integration can occur is through causal explanations for the updated information. Rapp and Kendeou (2009) found that people were better at updating their mental models for narratives when a characteristic (e.g., being clumsy) was followed by an inconsistency with a causal explanation (e.g., the person always had great balance *but the floor was just waxed*, emphasis added) rather than an inconsistency alone. They posit that a causal explanation provides a

stronger reason for people to update their understanding and integrate the new information with the old.

Belief rates in the present study may have been higher if the alternative explanations were causal. Experiment 3 materials were more causal than those in Experiments 1 and 2, in that the alternatives could be considered causes of the situation (e.g., contaminated equipment or pesticides causing spinach to become tainted, as compared to elderly adults or hockey players simply being passengers in an accident). However, to substantially increase belief and reduce the CIE, there may need to be a causal explanation for why the incorrect information was put forth originally. This is an interesting issue that should be further explored. In addition, the knowledge revision and refutation work of Kendeou, O'Brien, Rapp, and colleagues could also benefit from an exploration of the role of belief on corrections.

Another idea is that the CIE likely involves some degree of selective directed forgetting (e.g., Sahakyan, 2004) because people need to disregard a subset of information based on content. However, it is not clear the extent to which the processes involved in directed forgetting effects, such as differential rehearsal or suppression, are involved in CIE production when people do not accept the retraction. When it is accepted, the CIE is small, suggesting that any directed forgetting process is effective. Some evidence suggests that the fickle nature of selective directed forgetting effects may be due to how well integrated the to-be-forgotten information is with the rest (Pettijohn, 2016), which also aligns with the KReC framework (Kendeou et al., 2017; Kendeou & O'Brien, 2014). The current findings suggest that it may be of interest to explore whether people in selective directed forgetting studies accept the instruction to forget some items but not others.

### Limitations

A curious challenge in how we assess belief and acceptance is that it is possible that people who use retracted information may implicitly feel a need to say that they did not believe the retraction. Likewise, people who do not use retracted information could feel a need to say that they did believe the retraction. In other words, perhaps their earlier responses affected their retrospective assessment of retraction acceptance. This may or may not be the case, and it would be difficult, if not impossible, to disentangle these issues.

Similarly, it is difficult to determine in our study if participants accept the retraction (or not) while they are reading, or not until they are asked. We chose to place our belief assessment at the end to avoid inadvertently influencing people to process information differently while reading than they otherwise would, the issue of which is more relevant here than is the issue of acceptance timing. That said, Rapp and Kendeou (2009) found that people revised their understanding directly

after reading sentences refuting prior knowledge. It is possible that the same pattern occurs with retraction. More work is needed to determine what effect, if any, timing of the acceptance question has on the CIE.<sup>3</sup>

Generalizations about the present findings should be made cautiously. We did find support for the idea that belief influences the CIE across three experiments and three scenarios. However, Experiment 3 had a greater remaining CIE compared to the other experiments, despite a higher level of belief in the alternative condition. Belief did not have as strong an effect overall. It is possible that the within-participants design led people to respond differently across the scenarios than they otherwise would. Further research to examine the influence of design, as well as with different material types, is needed to determine the generality of the finding.

## Conclusion

In summary, in our study of the continued use of retracted information, we found that part of this finding may be due to a failure to accept the correction, leading people to continue using the original, inappropriate information. Work in this area needs to take this factor into account. The magnitude of the CIE may be exaggerated not by memory processes from a prior understanding, but from an unwillingness to accept the retraction.

Our findings may have real-world implications. We found that the majority of participants who were presented with a retraction did not accept it, even though it came from authority figures. This suggests that, in real-life situations, people may not accept retractions even from authority figures. Because news reports frequently are updated as more information becomes known, corrections and retractions are unavoidable. Therefore, to improve retraction acceptance, more work needs to be done on the way information is initially presented, as well as how retractions or corrections are presented.

**Author note** This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1313583, awarded to Andrea E. O’Rear. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

**Open practices statement** The data and materials for all experiments are available upon request. None of the experiments were preregistered.

<sup>3</sup> Ongoing work in the authors’ lab suggests that timing of when the retraction is presented may influence the CIE.

## Appendix 1

### Minibus accident (cf. Ecker et al., 2010)

#### Messages version

**Reading instructions (Exp. 3)** What you are to do in this part of the study is to carefully read six sets of messages. Try to read as you normally would. For each set, you will be shown one message at a time. After reading each message, you should click the NEXT button on the screen to advance to the next message. The title of each set will appear in red before the set begins. When you have finished reading all six sets of messages, you will be given instructions for the next part of the study. If you have any questions, please feel free to ask the experimenter.

1. A report has come in to the police headquarters on Sunday about a serious accident involving a minibus.
2. The report came from a driver who was driving past the scene of the accident.
3. The minibus had crashed into a steep embankment near Spring Street and had rolled on its side.
4. This resulted in injuries to some of the passengers on board.
5. A rescue crew was dispatched to the scene immediately upon report of the accident.
6. They arrived at the scene within ten minutes.
7. Police have stated that the passengers on the bus were a group of elderly people who were on their way back to their nursing home after a bingo game.
8. The weather was reportedly fine and visibility was good.
9. No other vehicles seem to have been involved.
10. The rescue crew started evacuation of the minibus and wanted to reach the injured passengers first but found it difficult to tell them apart from the uninjured.
11. The rescue crew also reported difficulty in getting both injured and uninjured passengers out of the minibus even though the exits were clear.
12. Authorities are desperately trying to reach the family members of the victims involved in the accident and managed to trace the minibus license number to a rental company.
13. The rescue crew reported that the rescue efforts were slow and would take more time and advised bystanders not to crowd around the area.
14. Due to building-up traffic, the local radio station advised drivers to avoid the Spring Street area.
15. Local television shows live footage of uninjured passengers having problems getting up the embankment.
16. Rescue crew can be heard remarking that the uninjured passengers were unable to help in the rescue efforts.

17. (control). A second statement from the police has confirmed that all passengers were successfully rescued.

17. (retraction). A second statement from the police stated that the passengers on the minibus were not elderly people.

17. (retraction and alternative). A second statement from the police has stated that the passengers on the minibus were not elderly people but college hockey players returning from a victory party after the state hockey final.

18. Some passengers with injuries were taken to the nearby St Joseph's hospital for treatment.

19. At the hospital, three passengers with more serious injuries had to be admitted for observation while the others were discharged after treatment.

20. Some of the uninjured passengers interviewed at the scene of the accident remarked how helpless and dependent they were on the rescue crew during the accident and expressed their gratitude to them.

#### Inference questions:

1. Why do you think it was hard telling apart who the injured and uninjured passengers were?
2. Why do you think it was difficult getting both the injured and uninjured passengers out of the minibus?
3. Which family members of passengers are authorities most likely to contact to inform them about the accident?
4. Why do you think it was difficult getting the uninjured passengers up the embankment?
5. Why do you think the uninjured passengers were unable to help with the rescue efforts?
6. Why do you think some passengers were injured while others were not?
7. How could such an incident be avoided in the future?
8. Why did the uninjured passengers feel helpless and dependent on the rescue crew?
9. Why do you think the minibus crashed?

#### Fact questions:

10. On which day did the accident occur?
11. Who reported the accident to the police?
12. Where did the minibus land after the accident?
13. Where did the accident occur?
14. How was the weather like on that day?
15. Which hospital were the injured taken to?
16. How many injured passengers were admitted to the hospital?
17. Where were the uninjured passengers interviewed?
18. Who was on the minibus during the accident?

#### Retraction memory:

19. What was the purpose of the second message from the police?
20. Were you aware of any corrections in the messages that you read?

#### Retraction belief:

21. Do you believe any corrections in the messages that you read, or did you think that they were wrong?

## Appendix 2

### Spinach recall

#### Narratives version

**Reading instructions (Exp. 3)** What you are to do in this part of the study is to carefully read six short stories. Try to read these stories as you normally would. For each story, you will be shown one sentence at a time. After reading each sentence, you should click the NEXT button on the screen to advance to the next sentence. The title of each story will appear in red before the story begins. When you have finished reading all six sets of messages, you will be given instructions for the next part of the study. If you have any questions, please feel free to ask the experimenter.

The news station received a report about a serious food recall in the Midwest. A large supply of baby spinach has been contaminated. Packages of all sizes of Green Acres brand fresh baby spinach are included in the recall. FDA officials have stated that the spinach was contaminated due to unsanitary equipment in the processing plant. Two dozen people have fallen ill after eating the spinach. The reported illnesses were located in Indiana, Illinois, and Missouri. All affected people were admitted to their local hospitals with severe gastrointestinal illness and dehydration. Four of the cases are reportedly highly severe, meaning that the patients have undetermined release dates from the hospital. Several employees complained about experiencing an unpleasant odor when they were in certain areas of the plant building. None of the employees formally filed any complaints. The FDA has ordered mandatory re-training in food safety for all Green Acres supervisors and shift workers. In an interview, the president of Green Acres expressed a sincere apology to all those affected by the contaminated spinach and promised full cooperation with the FDA investigation. The president assured the public that this is an isolated incident due to worker negligence and is not a reflection of the quality of Green Acres produce. (control) A second statement from the FDA reported that the



situation is under control and no more contaminated spinach will be shipped. // (retraction) A second statement from the FDA reported that the spinach was not contaminated due to unsanitary equipment. // (retraction + alternative) A second statement from the FDA reported that the spinach was not contaminated by unsanitary equipment, but by pesticide that had expired past its use date. All products packaged with the contaminated batch of spinach are being removed immediately from grocery stores. Consumers are urged to throw away all Green Acres baby spinach products purchased with a “sell by” date from June 7<sup>th</sup> to June 15<sup>th</sup>. If anyone has additional questions or concerns, they may call the Green Acres customer service at (800) 489-0501.

#### Inference questions:

1. What do you think made people sick after eating the spinach?
2. Why do you think the illness was severe enough to send people to the hospital?
3. What do you think the mandatory re-training is likely to involve?
4. What was the cause of the unpleasant odor?
5. Why do you think the employees did not file any complaints?
6. Why do you think the president blamed worker negligence?
7. What do you think the FDA investigation will focus on?
8. Why did the contamination affect such a large amount of spinach?
9. How could such an incident be avoided in the future?

#### Fact questions:

10. In what region of the country was the recall located?
11. What was the name of the brand of spinach?
12. How many people fell ill?
13. What symptoms did the sick people have?
14. How many cases were highly severe?
15. What three states were the sick people from?
16. Who did the FDA order to be re-trained?
17. Who issued an apology?
18. What was the month of the "sell by" dates?

#### Retraction memory:

19. What was the purpose of the second message from the FDA?
20. Were you aware of any corrections in the messages that you read?

#### Retraction belief:

21. Do you believe any corrections in the messages that you read, or did you think that they were wrong?

## References

- Akan, M., & Sahakyan, L. (2018). Repeated failures to obtain selective directed forgetting in lab and online samples and variations in stimuli. *Memory*, 26(3), 294-305.
- Amlund, J. T., Kardash, C. A. M., & Kulhavy, R. W. (1986). Repetitive reading and recall of expository text. *Reading Research Quarterly*, 21(1), 49-58. <https://doi.org/10.2307/747959>
- 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), 1-21.
- Delaney, P. F., Nghiem, K. N., & Waldum, E. R. (2009). The selective directed forgetting effect: Can people forget only part of a text? *The Quarterly Journal of Experimental Psychology*, 62(8), 1542-1550.
- Ecker, U. K., Lewandowsky, S., & Apai, J. (2011). Terrorists brought down the plane!—No, actually it was a technical fault: Processing corrections of emotive information. *The Quarterly Journal of Experimental Psychology*, 64(2), 283-310. <https://doi.org/10.1080/17470218.2010.497927>
- Ecker, U. K., Lewandowsky, S., Fenton, O., & Martin, K. (2014). Do people keep believing because they want to? Preexisting attitudes and the continued influence of misinformation. *Memory & Cognition*, 42(2), 292-304. <https://doi.org/10.3758/s13421-013-0358-x>
- Ecker, U. K., Lewandowsky, S., Swire, B., & Chang, D. (2011). Correcting false information in memory: Manipulating the strength of misinformation encoding and its retraction. *Psychonomic Bulletin & Review*, 18(3), 570-578. <https://doi.org/10.3758/s13423-011-0065-1>
- Ecker, U. K., Lewandowsky, S., & Tang, D. T. (2010). Explicit warnings reduce but do not eliminate the continued influence of misinformation. *Memory & Cognition*, 38(8), 1087-1100. <https://doi.org/10.3758/MC.38.8.1087>
- Fritz, C. O., Morris, P. E., Bjork, R. A., Gelman, R., & Wickens, T. D. (2000). When further learning fails: Stability and change following repeated presentation of text. *British Journal of Psychology*, 91(4), 493-511. <https://doi.org/10.1348/000712600161952>
- Greitemeyer, T. (2014). Article retracted, but the message lives on. *Psychonomic Bulletin & Review*, 21(2), 557-561. <https://doi.org/10.3758/s13423-013-0500-6>
- Guillory, J. J., & Geraci, L. (2010). The persistence of inferences in memory for younger and older adults: Remembering facts and believing inferences. *Psychonomic Bulletin & Review*, 17(1), 73-81. <https://doi.org/10.3758/PBR.17.1.73>
- 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(4), 201-209. <https://doi.org/10.1016/j.jarmac.2013.10.001>
- Hamm, V. P., & Hasher, L. (1992). Age and the availability of inferences. *Psychology and Aging*, 7(1), 56-64. <https://doi.org/10.1037/0882-7974.7.1.56>
- Howe, M. J. (1970). Repeated presentation and recall of meaningful prose. *Journal of Educational Psychology*, 61(3), 214-219. <https://doi.org/10.1037/h0029102>

- Howe, M. J. (1972). Repeated presentation and retention of meaningful information. *Psychological Reports*, 31(3), 840–842. <https://doi.org/10.2466/pr0.1972.31.3.840>
- 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(6), 1420–1436. <https://doi.org/10.1037/0278-7393.20.6.1420>
- Kang, S. H., Pashler, H., Cepeda, N. J., Rohrer, D., Carpenter, S. K., & Mozer, M. C. (2011). Does incorrect guessing impair fact learning? *Journal of Educational Psychology*, 103(1), 48–59. <https://doi.org/10.1037/a0021977>
- Kay, H. (1955). Learning and retaining verbal material. *British Journal of Psychology*, 46(2), 81–100. <https://doi.org/10.1111/j.2044-8295.1955.tb00527.x>
- Kendeou, P., Butterfuss, R., van Boekel, M., & O'Brien, E. J. (2017). Integrating relational reasoning and knowledge revision during reading. *Educational Psychology Review*, 29, 27–39. <https://doi.org/10.1007/s10648-016-9381-3>
- Kendeou, P., & O'Brien, E. J. (2014). The knowledge revision components (KReC) framework: processes and mechanisms. In D. N. Rapp & J. L. G. Braasch (Eds.), *Processing inaccurate information: Theoretical and applied perspectives from cognitive science and the educational sciences* (pp. 353–377). Cambridge: MIT Press.
- Kendeou, P., Smith, E. R., & O'Brien, E. J. (2013). Updating during reading comprehension: Why causality matters. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(3), 854–865.
- Lewandowsky, S., Ecker, U. K. H., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest*, 13(3), 106–131. <https://doi.org/10.1177/1529100612451018>
- Lewandowsky, S., Oberauer, K., & Gignac, G. E. (2013). NASA faked the moon landing—therefore, (climate) science is a hoax: An anatomy of the motivated rejection of science. *Psychological Science*, 24(5), 622–633. <https://doi.org/10.1177/0956797612457686>
- Lewandowsky, S., Stritzke, W. G. K., Oberauer, K., & Morales, M. (2005). Memory for fact, fiction, and misinformation: The Iraq War 2003. *Psychological Science*, 16(3), 190–195. <https://doi.org/10.1111/j.0956-7976.2005.00802.x>
- Malpass, R. S., & Devine, P. G. (1981). Eyewitness identification: Lineup instructions and the absence of the offender. *Journal of Applied Psychology*, 66, 482–489. <https://doi.org/10.1037/0021-9010.66.4.482>
- O'Brien, E. J., Cook, A. E., & Guéraud, S. (2010). Accessibility of outdated information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(4), 979–991.
- Pettijohn, K. A. (2016). Selective directed forgetting and event models. Unpublished Doctoral Dissertation.
- Prasad, M., Perrin, A. J., Bezila, K., Hoffman, S. G., Kindleberger, K., Manturuk, K., & Smith Powers, A. (2009). “There must be a reason”: Osama, Saddam, and inferred justification. *Sociological Inquiry*, 79(2), 142–162. <https://doi.org/10.1111/j.1475-682X.2009.00280.x>
- Radvansky, G. A., & Copeland, D. E. (2006). Walking through doorways causes forgetting: Situation models and experienced space. *Memory & Cognition*, 34(5), 1150–1156.
- Radvansky, G. A., & Copeland, D. E. (2010). Reading times and the detection of event shift processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(1), 210–216. doi:10.1037/a0017258
- Rapp, D. N., & Kendeou, P. (2007). Revising what readers know: Updating text representations during narrative comprehension. *Memory & Cognition*, 35(8), 2019–2032.
- Rapp, D. N., & Kendeou, P. (2009). Noticing and revising discrepancies as texts unfold. *Discourse Processes*, 46, 1–24. <https://doi.org/10.1080/01638530802629141>
- 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(1), 62–74. <https://doi.org/10.1037/xlm0000155>
- Sahakyan, L. (2004). Destructive effects of “forget” instructions. *Psychonomic Bulletin & Review*, 11(3), 555–559.
- Sparks, J. R., & Rapp, D. N. (2011). Readers’ reliance on source credibility in the service of comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(1), 230–247. <https://doi.org/10.1037/a0021331>
- Stebly, N. M. (1997). Social influence in eyewitness recall: A meta-analytic review of lineup instruction effects. *Law and Human Behavior*, 21, 283–297. <https://doi.org/10.1023/A:1024890732059>
- Storm, B. C., Koppel, R. H., & Wilson, B. M. (2013). Rapid communication: Selective cues to forget can fail to cause forgetting. *Quarterly Journal of Experimental Psychology*, 66(1), 29–36.
- Wilkes, A. L., & Leatherbarrow, M. (1988). Editing episodic memory following the identification of error. *The Quarterly Journal of Experimental Psychology*, 40(2), 361–387. <https://doi.org/10.1080/02724988843000168>
- Wilkes, A. L., & Reynolds, D. J. (1999). On certain limitations accompanying readers’ interpretations of corrections in episodic text. *The Quarterly Journal of Experimental Psychology: Section A*, 52(1), 165–183. <https://doi.org/10.1080/713755808>
- Zwaan, R. A. (1994). Effect of genre expectations on text comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(4), 920–933.
- Zwaan, R. A., Magliano, J. P., & Graesser, A. C. (1995). Dimensions of situation model construction in narrative comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(2), 386–397.
- Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, 123(2), 162–185.

**Publisher’s note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.