Journal of Experimental Psychology:
Learning, Memory, and Cognition
2016, Vol. 42, No. 1, 62–74

0278-7393/16/$12.00 http://dx.doi.org/10.1037/xlm0000155

The Continued Influence of Implied and Explicitly Stated Misinformation in News Reports

Patrick R. Rich and Maria S. Zaragoza
Kent State University

The piecemeal reporting of unfolding news events can lead to the reporting of mistaken information (or misinformation) about the cause of the newsworthy event, which later needs to be corrected. Studies of the *continued influence effect* have shown, however, that corrections are not entirely effective in reversing the effects of initial misinformation. Instead, participants continue to rely on the discredited misinformation when asked to draw inferences and make judgments about the news story. Most prior studies have employed misinformation that explicitly states the likely cause of an outcome. However, news stories do not always provide misinformation explicitly, but instead merely imply that something or someone might be the cause of an adverse outcome. Two experiments employing both direct and indirect measures of misinformation reliance were conducted to assess whether implied misinformation is more resistant to correction than explicitly stated misinformation. The results supported this prediction. Experiment 1 showed that corrections reduced misinformation reliance in both the explicit and implied conditions, but the correction was much less effective following implied misinformation. Experiment 2 showed that implied misinformation was more resistant to correction than explicit misinformation, even when the correction was paired with an alternative explanation. Finally, Experiment 3 showed that greater resistance to correction in the implied misinformation condition did not reflect greater disbelief in the correction. Potential reasons why implied misinformation is more difficult to correct than explicitly provided misinformation are discussed.

Keywords: continued influence, misinformation, memory, correction, judgments

In our information-rich society, people have constant access to the media through radio, television, and the Internet. The media uses these formats to report up-to-the-minute news about political figures, current events, and the state of the world. Although most of the news is factually correct, some news is based on incomplete information, speculation, or mistaken information. In cases where the news is reported in error, news organizations typically issue corrections and revise the story. Ideally, such corrections should be sufficient to counteract the initial error. To the contrary, laboratory studies of the *continued influence effect* have consistently demonstrated that when misinformation provides a causal explanation for a newsworthy outcome, people continue to rely on the discredited misinformation, even when they can remember and report the correction (Ecker, Lewandowsky, & Apai, 2011; Ecker, Lewandowsky, Fenton, & Martin, 2014; Ecker, Lewandowksy, Swire, & Chang, 2011; Ecker, Lewandowsky, & Tang, 2010; Guillory & Geraci, 2010, 2013; Johnson & Seifert, 1994, 1999; Seifert, 2002; Wilkes & Leatherbarrow, 1988; Wilkes & Reynolds, 1999; for recent reviews, see Lewandowsky, Ecker, Seifert, Schwarz, & Cook, 2012 and Ecker, Swire, & Lewandowsky, 2014).

In the experimental paradigm used to study the continued influence effect, participants read a series of messages imitating an unfolding news report in which a target piece of mistaken information (hereafter referred to as "misinformation") is later corrected.[1] For example, in a study by Johnson and Seifert (1994, Experiment 3b), participants read a news report describing the theft of valuable jewelry from a couple's home that occurred while they were away on vacation. The news story initially reports that the prime suspect in the case is the couple's son, who had been had been asked to check in on the house while the couple was away. Later in the same news story, this initial piece of misinformation is corrected, and participants learn that the son is no longer a suspect because he was out of town when the theft occurred. In studies of the continued influence effect, participants then receive a comprehension questionnaire consisting of open-ended prompts, factual questions, and inference questions. The primary question of

This article was published Online First July 6, 2015.

Patrick R. Rich and Maria S. Zaragoza, Department of Psychological Sciences, Kent State University.

Parts of the research reported in this article are based on a master's thesis conducted by Patrick R. Rich in partial fulfillment of the requirements for the MA degree, Department of Psychological Sciences, Kent State University. We thank Anna Aho-Pynttäri, Sean Burridge, Ashlee Chan, Meagan DeWitt, Bethany Dibble, Alexandra Evans, Jamie Garrod, Emily Grande, Katheryn Hrabik, Elizabeth Hyclak, Katelyn Lawrence, Kate Lenart, Jenn Lester, Meghan Mullett, Richard Phend, Melanie Rayk, Lauren Rothenbusch, Angelina Shirey, and Nicole Stevens for assistance with data collection, data entry, and data coding.

Correspondence concerning this article should be addressed to Patrick R. Rich, Department of Psychological Sciences, Kent State University, Kent, OH 44240. E-mail: prich1@kent.edu

[1] "Misinformation" is the term used in studies of the continued influence effect to refer to initial, mistaken information that is later corrected in the news story. We use the same term here to remain consistent with the prior literature.

interest is whether the correction is effective at undoing the effects of the misinformation as evidenced by performance on this questionnaire.

The dependent variable of primary interest in addressing this question is performance on a series of inference questions, which require participants to go beyond the events reported in the story in generating a response. For example, Johnson and Seifert (1994) asked participants the inference question, "How did the thief enter the house?" even though the story never mentioned how the thief entered. The measure of continued influence is the extent to which participants' rely on the corrected misinformation in responding to the inference questions, either by making direct references to the misinformation or by providing responses that are thematically consistent with the misinformation. For example, in Johnson and Seifert (1994, Experiment 3b) responses that reflected an "inside job" (e.g., "the thief used a key") were coded as consistent with the misinformation that the son stole the jewelry, and responses reflecting an "outside job" (e.g., "the thief entered through a broken window") were not. The continued influence effect is the finding that participants who receive misinformation that is later corrected make more inferences consistent with the discredited misinformation than control participants who were never misled.

The finding that participants continue to rely on the corrected misinformation when responding to inference questions is highly reliable and has been replicated across both a wide variety of news stories (e.g., a warehouse fire, a missing person, a terrorist attack, a bus accident, a house theft) and a range of different types of misinformation (Ecker, Lewandowsky, & Apai, 2011; Ecker et al., 2010; Johnson & Seifert, 1994; Wilkes & Leatherbarrow, 1988). These studies have shown that corrections reduce reliance on the misinformation but do not eliminate it. Studies have shown that corrections alone are largely ineffective at eliminating the effects of the initial misinformation, regardless of whether the correction is provided immediately after the misinformation (Johnson & Seifert, 1994), whether the correction is presented on multiple occasions (Ecker, Lewandowsky, Swire, & Chang, 2011), or whether people are given generic warnings about false information in news reports (Ecker et al., 2010). It is important that these findings cannot be attributed to poor memory for the correction, because the continued influence effect has been repeatedly demonstrated among participants who can remember and report the correction. As such, the empirical literature on the continued influence effect clearly demonstrates that the correction does not remove the misinformation, but rather that the misinformation and the correction coexist in memory.

A second factor that contributes to the continued influence effect is people's failure to monitor the validity of the misinformation when it comes to mind (Ecker et al., 2010; Ecker, Swire, & Lewandwosky, 2014). Such monitoring processes are a common component of frameworks and models of memory (Ayers & Reder, 1998; Gallo, 2010; Johnson, Hashtroudi, & Lindsay, 1993; Moscovitch & Melo, 1997; Schacter, Norman, & Koutstaal, 1998, to name a few) and refer to those strategic retrieval and monitoring processes involved in assessing the source, relevance, and validity of retrieved information. When responding to the inference questions, the retracted misinformation is likely to be activated automatically because it is closely linked to the elements of the story and provides a relevant basis for responding. Continued influence effects occur when participants fail to identify that the activated misinformation is invalid, either because they do not retrieve the correction or because they fail to engage in the systematic and effortful reasoning required for successful discrediting of the misinformation and generating a response on some other basis. These resource-demanding, controlled processes might involve, for example, noting the discrepancy between the misinformation and the subsequent correction, reasoning that the misinformation can no longer be true given the correction, concluding that the misinformation is not a valid basis for responding, withholding the misinformation as a response, revising one's initial beliefs in light of the correction, and generating a response based on the revised belief. In sum, people fall back on the discredited misinformation when responding to the inference questions because it is easy to do so.

Current Study

The impetus for the current study is the observation that news stories do not always provide misinformation explicitly, but instead merely imply that something or someone might be the cause of an adverse outcome. For example, Seifert (2002) describes an actual news story where a television news channel reported that a family of four had been found dead in their homes on the same evening that they had eaten at a Chinese restaurant in town. A few days later, the news channel reported that the deaths were actually caused by a faulty furnace. This story concluded with a report that the Chinese restaurant had closed, presumably because people could not shake the suspicion that something wrong with the restaurant and avoided dining there. In this story, the conclusion that the family died of food poisoning is implied by the story but never stated explicitly. As such, the inference that the deaths were caused by dining at the restaurant is an inference generated by the reader. In the experiments reported in this study, we tested the hypothesis that implied misinformation may be more resistant to correction than explicitly provided misinformation.

Implied misinformation may be especially difficult to correct for several reasons. First, whereas explicit misinformation provides a direct statement that the misinformation caused the outcome, implied misinformation requires that participants self-generate, or infer, the cause of the outcome implied by the news story. As a consequence, in the implied case, readers have to form connections between the evidence and the implied cause, and between the implied cause and the outcome resulting in a more elaborate representation (see Myers, Shinjo, & Duffy, 1987, and Mason & Just, 2004, for evidence consistent with this prediction). Because these self-generated causal links are likely to be more densely integrated with the story representation than are those that are explicitly provided, this may render implied misinformation more difficult to correct than explicitly provided information, because more revision is required in the former case. Empirical evidence consistent with the proposal that self-generated causal inferences are especially difficult to correct can be found in the related domain of belief perseverance. For example, a study by Davies (1997) showed that participants who (overtly) self-generated explanations for an event outcome evidenced more belief perseverance following an evidential discrediting manipulation than participants provided with an explanation for the outcome.

Second, when misinformation is explicitly provided, it may be relatively easy for readers to detect that a correction invalidates the

misinformation. In contrast, when misinformation is only implied, readers may be less able to detect the inconsistency between the correction and the misinformation. Hence, implied misinformation may be more difficult to correct than explicit misinformation because readers have more difficulty detecting this discrepancy. This prediction is consistent with a theory of mental contamination put forward by Wilson and Brekke (1994), who argue that an individual must be both aware of a mental contaminant to correct it and aware that the correction is relevant to utilize it. Support for the idea that participants may be less aware of implied relative to explicitly stated misinformation comes from research on the effects of innuendo in headlines. These studies have shown that the persuasive effects of incriminating innuendo were more difficult to counteract (e.g., by attributing them to a low credibility source) than the persuasive effects of directly stated incriminating evidence (Wegner, Wenzlaff, Kerker, & Beattie, 1981), presumably because participants are less aware of the persuasive effects of the innuendo (see also Ecker, Lewandowsky, Chang, & Pillai, 2014). In sum, for all of the aforementioned reasons, we hypothesized that implied misinformation would be more resistant to correction than explicitly provided misinformation.

To assess whether implied misinformation is more resistant to correction than explicit misinformation, we modified materials and procedures employed by Johnson and Seifert (1994, Experiment 3b). Following Johnson and Seifert, participants in the explicit misinformation conditions read a news story about a house theft where the identity of the presumed thief is explicitly stated (e.g., "Police suspect that the Harter's son, Evan, may have taken the box from the house"). Many of the details reported in the story were consistent with the idea that the owner's son committed the theft (e.g., the son watched over the house while the owners were away, the son had gambling debts, the son did odd jobs around the neighborhood, the neighborhood had experienced a rash of thefts). However, the story also reported details consistent with the alternative possibility that a random burglar committed the theft (e.g., the basement window was broken). To create an implied misinformation condition, participants read the identical story, but with the explicit statement about the son being a suspect removed. Because the preponderance of evidence provided in the story pointed to the son as the likely culprit, we assumed that participants in the implied misinformation condition would infer that the son committed the theft. For both the implied and explicit misinformation conditions, some participants later encountered the correction that the son had been out of town, and some did not.

A second innovation of the current study was that we employed more comprehensive measure of continued belief in corrected misinformation. In addition to the inference question measure that has been employed in all prior studies of continued influence, we used a second, more direct, measure of continued influence, where we asked participants to indicate to what extent they believed that the corrected misinformation was responsible for the outcome reported in the news story.[2] Specifically, we asked participants, "*How much do you believe that Evan Harter was involved in the theft of the jewelry?*" and they responded on a scale from 1 (*highly unlikely*) to 6 (*highly likely*). Whereas the inference question measure provides an indirect means of assessing continued reliance on the misinformation (by assessing the extent to which, when probed, participants generate inferences thematically consistent with the misinformation), reliance on the misinformation should also manifest as continued belief in the validity of the corrected misinformation. Our novel measure was designed to probe participants directly about their belief in the corrected misinformation. As we will show, performance on the inference question and direct measure were correlated, thus verifying that they are complementary measures of misinformation reliance.

In the studies reported here, we analyzed the shared variance between the inference question and direct measure with a multivariate analysis of variance (MANOVA), rather than analyze performance on each measure separately. Analyzing the two measures together with a MANOVA yields a cleaner, more robust and global measure of misinformation reliance that is independent of measure-specific characteristics. Accordingly, the primary variable of interest is "misinformation reliance" as assessed by this MANOVA.

The primary goal of the current study was to test the hypothesis that implied misinformation is more resistant to correction than explicit misinformation. To test this prediction, we compared each of the corrected misinformation conditions (i.e., corrected implied and corrected explicit) to their corresponding uncorrected misinformation conditions (uncorrected implied and uncorrected explicit). Given the hypothesis that implied misinformation is more difficult to correct, we predicted that the magnitude of this correction effect (i.e., the extent to which the correction reduced the influence of the misinformation) would be smaller in the implied condition than in the explicit conditions.

Ideally, a correction should not only reduce reliance on the misinformation, but it should eliminate this reliance altogether. In other words, participants who read a correction should behave like those who have never been misinformed. Hence, the second method we used to gauge the effectiveness of the correction following implied and explicit misinformation was to assess whether misinformation reliance in the corrected misinformed conditions (implied and explicit) exceeded misinformation reliance in a baseline control condition, where participants had not been misled. This comparison served as the measure of the continued influence effect. We predicted that misinformation reliance following a correction would be greater in the implied condition than in the explicit condition.

Experiment 1

Method

Participants. A total of 357 undergraduates at a large Midwestern university (ages 18–28, $M = 19.41$; 211 women) completed the experiment and were awarded credit for a course requirement. All participants were fluent in English.

Design. The design of Experiment 1 consisted of a 3 (Misinformation Type: control, implied, or explicit) $\times$ 2 (Correction Condition: uncorrected or corrected) between-subjects factorial design. An a priori power analysis (G*Power 3; Faul, Erdfelder,

[2] This is not to say that prior studies have employed only the inference question measure. To the contrary, prior studies have routinely included other free response measures (e.g., free recall of the news story and questions about the cause of the outcome). However, because participants rarely produce the misinformation on these measures, they have not been very diagnostic in assessing continued influence.

Lang, & Buchner, 2007) using the medium effect size for the continued influence effect found in Johnson ad Seifert (1994, Experiment 3b; $d = .59$, $\alpha = .05$, $1\text{-}\beta = .85$) indicated that 42 subjects should be collected in each condition. We collected approximately 18 additional subjects in each condition to compensate for the possible need to exclude participants from the study (for reasons described later).

Materials and procedure. Participants came to the lab in groups of 1 to 10 for a single 45-min session. All participants were informed that they would be reading a series of messages and then asked questions about the events described in the messages.

Phase 1: News story. Participants were given 5 min to read a news report adapted from an experiment conducted by Johnson and Seifert (1994, Experiment 3b). The news report described a police investigation into a jewelry theft at a private home while the owners were on vacation. The report consisted of a series of 13 messages each presented on a single sheet, one-at-a-time, and ranging from one to four sentences. Three of the 13 messages (messages 4, 5, and 11) were *critical messages* that varied depending upon condition. The remaining messages were identical across all conditions (see Appendix A for the full list of messages).

The manipulation of misinformation type (explicit, implied, or control) was implemented in Message 5 (provided later). For participants in the explicit and implied conditions (both corrected and uncorrected), Message 5 contained three pieces of evidence implicating the owner's son as the thief, because the evidence suggested that the son had both opportunity and a potential motive for committing the theft. However, only participants in the explicit conditions received the initial, explicit statement about the son being a suspect (see underlined statement below). For participants in the Implied conditions the underlined statement naming the son as a suspect was removed. Message 5 for the explicit and implied conditions read:

> *Police suspect that the Harters' son, Evan, may have taken the box from the house. The Harters report that they had asked their son to check in on the house periodically during their absence. The son also did other odd jobs for many of the neighbors to pay off his recent gambling debts.*

For participants in the control condition, the news story mentioned that the owners had a son, but did not implicate him as the thief. To accomplish this, the information contained in Messages 4 and 5 was rearranged and modified to remove all information implicating the son as the thief. Messages 4 and 5 for the control condition read:

> *Message 4: Mrs. Harter specifically remembers that she went through her jewelry before leaving to get her pearl earrings. They were the last things she packed before the Harters' son, Evan, arrived to take them to the airport.*
>
> *Message 5: A tall tree arches near the bedroom window, but police have found no evidence of tampering with the window.*

For participants in all three corrected conditions, a correction message was added to their news story and appeared as Message 11. Note that for the corrected control condition, the underlined portion was deleted (because the control story did not mention the son watching the house); participants in the corrected implied and explicit conditions received the identical correction. Message 11 in corrected conditions read:

> *Police have found that Evan had been called away on business and had not been in town to look after the house during the Harters' vacation.*

This is a correction because the son could not have committed the theft if he was out of town. Participants in the uncorrected conditions (explicit, implied, and control) did not receive the correction message.

Phase 2: Filler questions. Participants were given 20 min to complete a set of questionnaires as filler tasks.

Phase 3: Final questionnaire. Participants were given 20 min to complete a written questionnaire that was a modified version of the measures used by Johnson and Seifert (1994). The questionnaire consisted of several subgroups of questions that appeared in the following order: one question about the cause of the crime, 9 factual questions, 10 inference questions, one direct question about participants' belief in the son's involvement, a free recall of the events, and two manipulation check questions. The various question types are described here and are reported in full in Appendix B.

Cause question. The first question asked participants what caused the box to be missing.[3]

Factual questions. The following nine questions (Questions 2–10) queried participants about the main events of the news story.

Inference questions. The following 10 questions (Questions 11–20) were inference questions that asked participants about information not explicitly mentioned in the news story. Nine of the inference questions were taken from Johnson and Seifert (1994), and one question (our question 18) is a new question substituted for one of the original questions. In prior studies, this is the measure that has traditionally been used to assess continued influence of misinformation following correction.

Direct measure of belief in son's involvement. Question 21, designed for the purposes of this study, probed participants directly about their belief in the corrected misinformation. Participants responded to the question, "How much do you believe that Evan Harter was involved in the theft of the jewelry?" using a scale from 1 (*highly unlikely*) to 6 (*highly likely*).

Free recall. Participants were then asked to freely recall everything they could remember about the news report.[4]

Manipulation check questions. The final two questions (23 and 24) were designed to assess participants' memory for the correction ("What did the story report about where Evan Harter was during the Harters' vacation?" and "What facts about the case did the police change their minds about, based on information they discovered later?").

[3] This question showed no evidence of continued influence: participants who received a correction rarely spontaneously mentioned that Evan was responsible for the theft (corrected explicit: 2%; corrected implied: 2%). Because Experiment 2 yielded similar results, this measure will not be discussed further.

[4] This measure showed no evidence of continued influence: no participants in the corrected conditions spontaneously reported that Evan was responsible for the theft. Because Experiment 2 yielded similar results, this measure will not be discussed further.

Inference question coding. Following Johnson and Seifert (1994), the inference questions were coded for responses consistent with an "inside job" theme, which included references to the son's potential involvement, having access to the house, intimate knowledge of the Harter's and their home, and so forth. Two trained raters, blind to condition, independently coded each response with a score of "1" if an "inside job" inference was provided and a score of "0" if not. Overall, rater agreement was 97%. Disagreements were resolved through discussion with a third, trained rater.

Results and Discussion

The manipulation check questions revealed that the percentage of participants who failed to remember the correction in the corrected explicit (21%) and corrected implied (12%) conditions were not reliably different, $\chi^2(1) = 1.71$, $p = .19$, $\varphi = .12$. Nevertheless, because it was important to ensure that any group differences in continued misinformation reliance could not be attributed to differences in memory for the correction, we excluded from the analyses reported below all participants who failed to report the correction.[5] After removing these participants, 328 remained: uncorrected control ($n = 62$), corrected control ($n = 51$), uncorrected explicit ($n = 59$), corrected explicit ($n = 49$), uncorrected implied ($n = 55$), and corrected implied ($n = 52$).

Preliminary analyses also confirmed that, as expected, performance on the direct measure and the inference question measure (Figure 1) was highly correlated ($r = .48$, $p < .001$). As discussed earlier, both measures of misinformation reliance were analyzed together using a 3 (Misinformation Type: control, explicit, or implied) $\times$ 2 (Correction: uncorrected or corrected) MANOVA, in an effort to obtain a cleaner, more robust measure of misinformation reliance. For this reason, the omnibus MANOVA was followed by planned contrasts, rather than univariate analyses of the individual measures (although, as noted in footnotes 6 and 7, the pattern of results was the same when the measures are considered separately).

The predicted interaction was highly reliable, $V = 0.06$, $F(4, 644) = 5.21$, $p < .001$, $\eta_p^2 = .03$, as were the main effects of correction, $V = 0.13$, $F(2, 321) = 23.34$, $p < .001$, $\eta_p^2 = .13$, and misinformation type, $V = 0.18$, $F(4, 644) = 15.68$, $p < .001$, $\eta_p^2 = .09$. To assess the locus of the interaction, we conducted a series of planned contrasts that are summarized in Table 1.

We first compared performance in the uncorrected conditions. In the absence of a correction, control participants (who were not misinformed) evidenced less belief in the misinformation than both the explicit (Contrast 1) and implied (Contrast 2) participants, who in turn did not differ from each other (Contrast 3). In sum, in the absence of a correction, exposure to implied and explicit misinformation led to comparable levels of belief in the misinformation.

The primary test of the study hypothesis involved assessing performance following the correction. Comparison of performance in the corrected conditions revealed that the implied condition evidenced significantly higher misinformation reliance than both the control (contrast 5) and explicit conditions (Contrast 6), but the explicit and control conditions did not differ reliably (Contrast 4). In sum, the results supported the prediction that continued misinformation reliance was greater in the implied condition than in the explicit condition.[6]

Finally, a comparison of performance in the corrected and uncorrected conditions was carried out to assess the magnitude of the correction effect for each misinformation type. The correction did not reliably affect performance in the control condition (Contrast 7), but reliably reduced misinformation reliance in both the explicit (Contrast 8) and implied misinformation (Contrast 9) conditions. A follow up 2 $\times$ 2 MANOVA analyzing only the explicit and implied conditions verified that the magnitude of the correction effect was greater in the explicit condition than in the implied condition, as evidenced by a significant interaction $V = 0.06$, $F(4, 644) = 5.21$, $p < .001$, $\eta_p^2 = .03$.

In summary, the hypothesis that implied misinformation would be more difficult to correct than explicit misinformation was supported. Although the correction reduced misinformation reliance in both the explicit and implied conditions, the correction was more effective following explicit misinformation. In addition, participants in the implied condition evidenced greater continued influence of corrected misinformation than did those in the explicit condition.

Although the absence of a reliable continued influence effect in the Explicit condition (i.e., misinformation reliance in the corrected explicit condition did not exceed corrected control) may seem inconsistent with prior findings, we note that misinformation reliance in both the corrected control and corrected explicit conditions was well above zero (Figure 1). This raises the possibility that the news story employed in the control condition may have, to a small extent, implicated the son as a potential suspect. Even though control participants never received information implicating the owners' son of wrongdoing (i.e., misinformation), the news story did mention that the owners' had a son who drove them to the airport. Given Gricean principles of relevance in communication (Grice, 1975), the fact that the owners' son was mentioned in the news story may have been interpreted by some as implicating the son.

Whether or not the control condition employed here represents a true no-misinformation control, there was clear evidence of a reliable continued influence effect in the implied condition. More important, the results supported the hypothesis that implied misinformation was more resistant to correction than explicit misinformation. Experiment 2 was designed to replicate these findings using a stronger correction.

Experiment 2

One technique that consistently and substantially reduces continued influence is pairing the correction of the misinformation with an alternative explanation for the outcome (e.g., when par-

[5] The pattern of results for Experiments 1, 2, and 3 were the same even when including participants who did not remember the correction.

[6] When separate univariate analyses were conducted on each of the measures individually, the same finding of reliable continued influence following implied, but not explicit, misinformation was observed in both measures. Following the correction, the implied condition evidenced reliably greater reliance on the misinformation than the control condition, direct measure: $t(101) = 3.22$, $p < .01$; inference measure: $t(101) = 2.96$, $p < .01$, but the explicit condition did not, direct measure: $t(98) = -1.00$, $p = .32$; inference measure: $t(98) = 1.41$, $p = .16$.

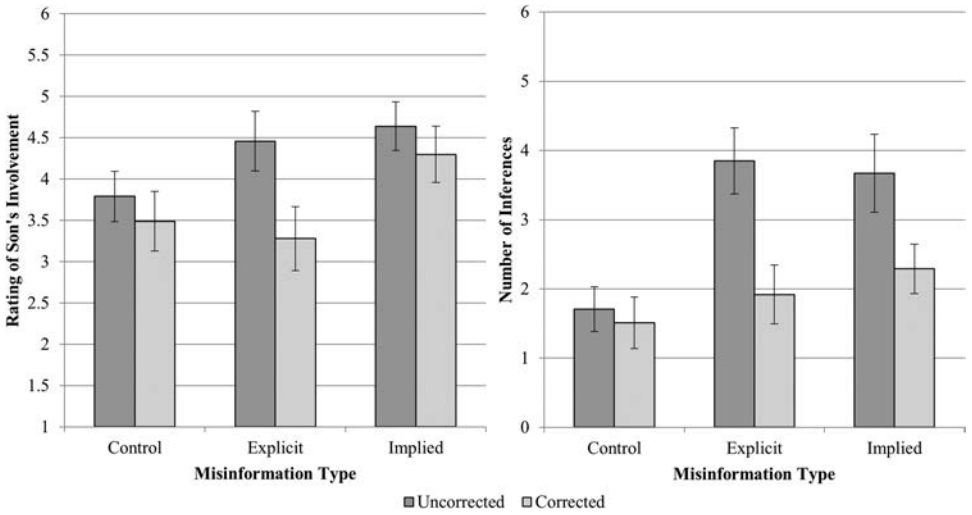

Figure 1. Effect of correction on misinformation reliance as a function of misinformation type for Experiment 1 (left panel: direct measure; right panel: inference measure). Bars represent 95% confidence intervals.

ticipants are informed that the actual thief was caught with the jewelry in his possession; Johnson & Seifert, 1994). This latter finding is consistent with the view that corrections are ineffective because a correction alone leaves a causal gap in the reader's mental model of the incident (Lewandowsky et al., 2012; Wilkes & Leatherbarrow, 1988). That is, a correction alone informs the reader that the misinformation is not responsible for the outcome, but it does not tell the reader the actual cause of the outcome. Thus, when readers have to draw inferences about the event or are asked what they believe caused the outcome, they may fall back on the misinformation because it is the only explanation they have. Pairing the correction of the misinformation with an alternative explanation remedies the situation by filling the causal gap, and provides the participant with a new, causally coherent account of the news event. Hence, the goal of Experiment 2 was to investigate whether implied misinformation would be more difficult to correct than explicit misinformation even in the face of an alternative explanation. Otherwise, Experiment 2 was identical to Experiment 1, with the exception of minor modifications outlined here.

Table 1
Planned Contrasts for Multivariate Analysis of Variance With "Inside Job" Inferences and Belief in the Son's Involvement as Outcome Variables, Experiment 1

Planned contrasts	Multivariate test results		
	$F(2, 321)$	p	η_p^2
Uncorrected conditions			
1. Uncorrected control vs. uncorrected explicit	26.29	<.001*	.14
2. Uncorrected control vs. uncorrected implied	22.13	<.001*	.12
3. Uncorrected explicit vs. uncorrected implied	0.91	.41	.01
Corrected conditions			
4. Corrected control vs. corrected explicit	2.18	.12	.01
5. Corrected control vs. corrected implied	6.02	.003*	.04
6. Corrected explicit vs. corrected implied	9.14	<.001*	.05
Correction effects			
7. Uncorrected control vs. corrected control	0.78	.46	.01
8. Uncorrected explicit vs. corrected explicit	22.79	<.001*	.12
9. Uncorrected implied vs. corrected implied	9.79	<.001*	.07

* Significant, Holm-Bonferroni correction.

Method

Participants. A total of 338 participants completed the experiment (ages 16–68, $M = 29.88$; 149 women, 50 unreported). An a priori power analysis (G*Power 3; Faul et al., 2007) using the effect size for the 2 × 2 MANOVA interaction from Experiment 1, $f^2(V) = .03$, $\alpha = .05$, 1-$\beta = .85$, suggested that 55 subjects should be collected in each condition. Of these, 106 participants were recruited from a large Midwestern university and were awarded credit for a course requirement. An additional 232 participants were recruited via Amazon's Mechanical Turk online marketplace. These participants were all located in the United States, had completed and been approved on at least 100 other tasks on Mechanical Turk, and at least 95% of their completed tasks had been approved. These participants completed the experiment through the Qualtrics online survey platform and were paid $1.50 for participating in the study. The pattern of results reported below remained the same when participants collected through Amazon's Mechanical Turk were analyzed separately. Accordingly, participants from both recruitment sources were combined. All participants were fluent in English.

The design, materials, and procedures were identical to Experiment 1, with the following exceptions. First, for participants in the corrected conditions (corrected implied, corrected explicit, and corrected control), the correction message (Message 11) was modified, such that the following alternative explanation for the outcome (based on Johnson & Seifert, 1994) came immediately after the correction that the son was out of town:

> *They have identified a suspect, ex-convict Dan Fowler. Fowler recently sold expensive jewelry to a pawn shop, and his girlfriend, Sarah, works for a cleaning service that regularly cleans the Harter's home.* (see Appendix C)

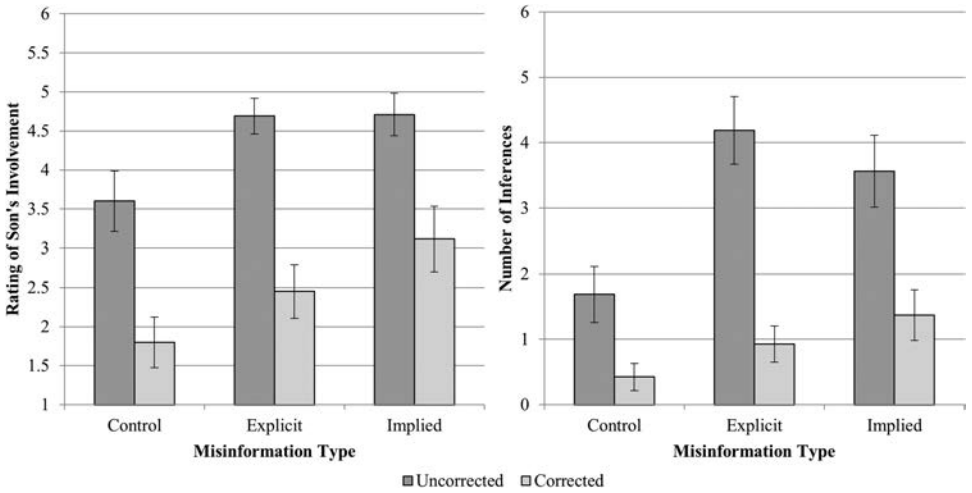

Figure 2. Effect of correction on misinformation reliance as a function of misinformation type for Experiment 2 (left panel: direct measure; right panel: inference measure). Bars represent 95% confidence intervals.

To reduce the time required to complete the study, we also removed the filler questions and the free-recall task.

Inference question coding. The coding procedure in Experiment 2 was identical to the procedure used in Experiment 1. Given that the coding scheme was highly reliable in Experiment 1, only 10% of the data were coded by a second coder. Rater agreement was 93%, and disagreements were resolved by a trained rater.

Results and Discussion

As in Experiment 1, the manipulation check questions revealed that the percentage of participants who failed to report the correction did not differ in the explicit and implied conditions (*M*s = 10% and 11%, respectively), $\chi^2(1) = 0.02$, $p = .88$, $\varphi = .01$. Nevertheless, we once again excluded participants who failed to remember the correction from the analyses reported below. After excluding these participants, a total of 316 remained: uncorrected control ($n = 48$), corrected control ($n = 40$), uncorrected implied ($n = 60$), corrected implied ($n = 56$), uncorrected explicit ($n = 61$), and corrected explicit ($n = 51$).

Preliminary analyses confirmed that performance on the direct measure and the inference measure (Figure 2) were again highly correlated ($r = .67$, $p < .001$). The measures of misinformation reliance were again analyzed using a 3 (Misinformation Type: control, explicit, or implied) × 2 (Correction: uncorrected or corrected) MANOVA.

The pattern of results replicated those of Experiment 1, with the exception that the correction reduced misinformation reliance even in the control condition. The predicted interaction was highly reliable, $V = 0.07$, $F(4, 620) = 5.32$, $p < .001$, $\eta_p^2 = .03$, as were the main effects of correction, $V = 0.43$, $F(2, 309) = 114.78$, $p < .001$, $\eta_p^2 = .43$, and misinformation Type, $V = 0.20$, $F(4, 620) = 16.99$, $p < .001$, $\eta_p^2 = .10$. Planned contrasts are reported in Table 2.

As in Experiment 1, comparisons of the uncorrected conditions revealed that misinformation reliance was lower in the control condition than in both the explicit (Contrast 1) and implied (Contrast 2) conditions, which in turn did not differ from each other (Contrast 3). Of primary relevance to the main hypotheses, the results revealed greater continued influence of corrected misinformation in the implied condition. Comparison of performance in the corrected conditions revealed that misinformation reliance in the implied condition exceeded that of both the control (Contrast 5) and explicit (Contrast 6) conditions, although the latter difference was only marginally significant when using a Holm-Bonferroni correction ($p_{corrected} = .06$). Misinformation reliance was also greater in the explicit condition relative to the control condition (Contrast 4), but, again, this difference was only marginally significant when using a Holm-Bonferroni correction ($p_{corrected} = .07$). Collectively, the foregoing analyses evidenced a pattern consistent with the finding in Experiment 1 that continued misinformation reliance was greater following implied, relative to explicit, misinformation.

Finally, a comparison of the uncorrected and corrected conditions was carried out to assess whether the magnitude of the correction effect varied as a function of misinformation type. Although there was a reliable correction effect in all three conditions (including the control condition, see Contrasts 7, 8, and 9), a follow up 2 × 2 MANOVA confirmed that the correction effect was larger in the explicit condition than in the implied condition, as evidenced by a significant interaction, $V = 0.03$, $F(2, 223) = 3.34$, $p = .04$, $\eta_p^2 = .03$. Hence, the results provide clear evidence that implied misinformation was more resistant to correction than explicit misinformation.[7]

The results of Experiment 2 replicate and extend the findings from Experiment 1 by showing that the pattern of results replicated even when participants were given a much stronger correction. In Experi-

[7] Again, the pattern of results remained the same when univariate analyses were conducted on each of the measures separately. Following a correction, the implied condition evidenced greater reliance on the misinformation than the explicit condition, direct measure: $t(105) = 2.45$, $p = .02$; inference measure: $t(105) = 1.85$, $p = .07$. In addition, both measures evidenced significant 2 (Correction) × 2 (Explicit vs. Implied) interactions, thus verifying that the correction effect was larger in the explicit condition than in the implied condition. Direct measure: $F(1, 223) = 3.97$, $p = .05$; inference measure: $F(1, 223) = 4.90$, $p = .03$.

Table 2
Planned Contrasts for Multivariate Analysis of Variance With "Inside Job" Inferences and Belief in the Son's Involvement as Outcome Variables, Experiment 2

Planned contrasts	Multivariate test results		
	$F(2, 309)$	p	η_p^2
Uncorrected conditions			
1. Uncorrected control vs. uncorrected explicit	31.78	<.001*	.17
2. Uncorrected control vs. uncorrected implied	20.55	<.001*	.12
3. Uncorrected explicit vs. uncorrected implied	2.80	.06	.02
Corrected conditions			
4. Corrected control vs. corrected Explicit	3.35	.04	.02
5. Corrected control vs. corrected implied	13.22	<.001*	.08
6. Corrected explicit vs. corrected implied	4.05	.02	.03
Correction effects			
7. Uncorrected control vs. corrected control	24.04	<.001*	.14
8. Uncorrected explicit vs. corrected explicit	73.70	<.001*	.32
9. Uncorrected implied vs. corrected implied	34.03	<.001*	.18

* Significant, Holm-Bonferroni correction.

ment 2, a correction accompanied by an alternative explanation for the outcome led to robust reductions in misinformation reliance, even in the control condition. Nevertheless, even when given this highly effective correction, implied misinformation evidenced more resistance to correction than explicit misinformation. Relative to participants in the explicit misinformation condition, participants in the implied misinformation condition evidenced greater reliance on the misinformation following the correction, and a smaller correction effect. In the next experiment, we assess whether the greater resistance to correction in the implied misinformation conditions reflects less willingness to accept the correction as true.

Experiment 3

In a postexperimental questionnaire administered after Experiment 1, we asked participants to what extent they believed the correction that the son was out of town when the theft occurred. Not only was the proportion of participants who endorsed believing the correction low, but we also found that participants in the implied condition evidenced lower belief in the correction (only 37% endorsed belief in the correction) than those in the explicit condition (where 47% endorsed believing the correction), though this difference was not statistically reliable, $\chi^2(1) = 1.12$, $p = .29$, $\varphi = .11$. Of course, given that participants had already endorsed their belief in the son's guilt on the study measures, their responses on the postexperimental questionnaire may have simply reflected an attempt to remain consistent with their prior responses. Nevertheless, this finding raised the possibility that participants in the implied condition may have been less willing to believe the correction than those in the explicit condition. By this account, participants in these conditions may have remembered the correction statement to the same extent (e.g., when answering the manipulation check questions), but those in the Implied condition may have been more likely to reject it as untrue or explain it away as not relevant (Otero & Kintsch, 1992). Rejecting the correction as untrue or irrelevant would make it unnecessary for participants to revise their belief in the misinformation. The goal of Experiment 3 was to investigate this potential mechanism by directly assessing to what extent participants in the implied and explicit conditions believed that the correction was true. To this end, in Experiment 3 we gave participants the news story to read, but removed from the questionnaire the two measures of misinformation reliance (direct measure and inference questions). Instead, we asked them to rate their belief in the correction.

Method

Participants. A total of 169 undergraduate participants from a large Midwestern university (age 18–43, $M = 19.93$; 128 women) completed the experiment and were awarded credit for a course requirement. All participants were fluent in English.

Design. The design was identical to Experiment 1 except that the uncorrected conditions were removed, leaving 3 between-subjects conditions: corrected control, corrected explicit, or corrected implied. An a priori power analysis (G*Power 3; Faul et al., 2007) using an assumed medium effect size for the omnibus analysis of variance (ANOVA) comparing all 3 conditions ($f = .25$, $\alpha = .05$, $1\text{-}\beta = .85$) suggested that 60 subjects should be collected in each condition.

Materials and procedure. The materials and procedure were identical to Experiment 1 with exception that the inference questions and direct question about the son's involvement were removed, and a measure of belief in the correction was added. Immediately after the manipulation check questions. participants were asked, "How much do you believe that Evan Harter was out of town on business while his parents were on vacation?" using a scale from 1 (*highly unlikely*) to 6 (*highly likely*).

Results and Discussion

As in Experiments 1 and 2, the rates of failing to report the correction on the manipulation check questions did not differ between the Explicit and Implied conditions, (4% and 3%, respectively, $\chi^2(1) = 0.001$, $p = .97$, $\varphi = .003$. After excluding participants who failed to report the correction, a total of 159 remained: corrected control ($n = 49$), corrected explicit ($n = 54$), and corrected implied ($n = 56$).

The results provided no support for the hypothesis that participants in the implied condition were more likely to disbelieve the correction than participants in the explicit misinformation condition. Mean ratings of belief that the son was out of town were 3.73, 3.31, and 3.23 for the corrected control, corrected explicit, and corrected implied conditions. A one-way ANOVA verified that there were no group differences, $F(2, 156) = 2.07$, $p = .13$, $\eta_p^2 = .03$. In sum, the results of Experiment 3 showed that the extent to which participants endorsed belief in the correction that the son was out of town did not differ for the explicit and implied conditions. Hence, the data provide no support for the hypothesis that implied misinformation is more resistant to correction because participants are more likely to reject the correction as untrue.

General Discussion

Misinformation presented in news stories can influence people's belief, impressions, and reasoning about the cause of a newsworthy event, and this influence persists even after people encounter

new information and evidence (i.e., corrections) that refutes the initial, mistaken information(Lewandowsky et al., 2012). The primary contribution of the current study was the novel finding that corrections are less effective when the misinformation is merely implied as opposed to explicitly provided. Although prior studies have examined the consequences of subtly misleading information in headlines (Ecker, Lewandowsky, Chang, et al., 2014; Wegner et al., 1981), this is the first study to our knowledge to directly compare the consequences of implied versus explicitly provided misinformation in news stories. In two experiments that employed both direct and indirect measures of misinformation reliance, we found clear and consistent evidence that implied misinformation was more resistant to correction than explicitly provided information. Implied misinformation evidenced greater resistance to correction even when the correction was accompanied by an alternative explanation for the outcome (Experiment 2), a manipulation that has been shown to greatly reduce and even eliminate the continued influence of misinformation (Ecker et al., 2010; Johnson & Seifert, 1994).

The findings reported here were obtained with a single news story that contained one piece of misinformation. Hence, the current findings will need to be replicated and extended to determine if they generalize to other news stories (e.g., current events, natural disasters, politics, more serious crimes, etc.), and other methods of implication (e.g., causal relationships implied by contiguous events; denials implying validity of the denied information). Given the evidence that continued influence effects are influenced by people's prior knowledge and beliefs (Ecker, Lewandowsky, Fenton, et al., 2014), additional research should assess whether the differences between implied and explicit misinformation interact with story content. For example, news stories that imply wrong doing by a person might be especially resistant to correction relative to news stories that imply a physical cause of an outcome (cf., Hilton, McClure, & Sutton, 2010). Clearly, there is much to be learned about the potential boundary conditions of the effects reported here. We note, however, that in the current experiments, implied misinformation was found to be more resistant to correction than explicit misinformation using two dependent variables (the inference question measure that is the standard in studies of continued influence, as well as a new, direct, measure of continued influence that we developed for purposes of this study). More important, our findings occurred across two correction types (Experiments 1 and 2) that varied considerably in their effectiveness. The consistency of our findings across these variables provides some confidence that they will generalize to other situations.

Why was implied misinformation more difficult to correct than explicitly provided misinformation? The findings from the present studies rule out several possibilities. First, we found no evidence that implied misinformation led to greater levels of initial belief in the misinformation prior to correction (in both experiments, performance in the uncorrected implied and explicit conditions did not differ). Hence, the differences in misinformation reliance observed after the correction cannot be attributed to differential levels of initial belief in implied versus explicit misinformation prior to receiving the correction. Second, in both experiments, we found no evidence that participants in the implied condition were less able to remember the correction. Furthermore, to ensure that any differences between implied and explicit misinformation performance could not be attributed to differential memory for the correction, the results reported here were conditional on memory for the correction. Hence, given that participants in the implied and explicit misinformation conditions of the current study started out with comparable levels of initial belief in the misinformation and remembered the correction equally well, we conclude that participants in the implied misinformation condition were likely either less able and/or willing to use the correction to revise their initial beliefs.

One way in which resistance to revising one's beliefs might be expressed is by simply rejecting the correction as untrue (cf., Otero & Kintsch, 1992). By this account, participants can remember the information contained in the correction when asked about it (e.g., they remember the news story reported that the son was out of town), but they reject the validity of this correction. Consistent with this idea, a number of studies have demonstrated that corrections contrary to a prior belief or worldview either have reduced the effect or sometimes even backfire (Hart & Nisbet, 2012; Lewandowsky, Stritzke, Oberauer, & Morales, 2005; Nyhan & Reifler, 2010; Nyhan, Reifler, & Ubel, 2013; for contrasting evidence see Ecker, Lewandowsky, Fenton, et al., 2014). Experiment 3 tested the possibility that participants would be less likely to believe the correction following implied, relative to explicit, misinformation, but found no evidence to support this possibility. Hence, it appears that participants in the implied and explicit misinformation conditions of the current study believed the correction to the same extent, but participants in the implied condition were less likely to use it.

Several cognitive mechanisms may alone, or in combination, underlie the finding that implied misinformation was more resistant to correction than explicit misinformation. First, with explicit misinformation the presumed cause of the outcome was explicitly provided, but with implied misinformation readers had to go beyond the evidence provided in the news story to infer, or self-generate, the likely cause of the outcome. Prior research has established that causal reasoning plays a central role in story comprehension, as the goal of comprehension is to identify the causal and logical sequence of events that lead to the outcome (see McNamara & Magliano, 2009), for a review. When the causal relations between the elements of a story are not explicitly provided (as with implied misinformation), the reader must generate inferences about the causal relationships between the elements of the story and integrate these inferences with the information provided in the story (e.g., inferring that the son may be committing the other thefts in the neighborhood when working his odd jobs; see Duffy, Shinjo, & Myers, 1990, and Kim & Van Dusen, 1998). The act of having to generate inferences may lead readers to activate relevant general knowledge schemas and scripts (e.g., about thefts, gambling debts, ne'er do well sons) that can serve as a basis for generating these inferences, but whose content may also become integrated with the story. On the other hand, when the causal relations between the elements of a story *are* explicitly provided (as with the explicit misinformation) the reader does not need to generate as many inferences to understand the causal relationship between the details in the story. Thus, because implied misinformation requires generating and integrating these causal inferences to a greater extent than explicit misinformation, the implied misinformation was likely integrated into a story representation that was richer and more highly interconnected than the story representation that resulted from reading the explicit news story.

We posit that the richer, more elaborate story representation generated by participants in the implied misinformation condition may have been more difficult to correct for several reasons. One

possibility is that, as a consequence of having self-generated the misinformation, participants in the implied condition were much more convinced about the validity of the misinformation than were participants in the explicit condition (who were only passively exposed to the misinformation). As a consequence, participants in the implied condition may have been less willing to revise their initial, strongly held, beliefs. Although the current findings cannot address this possibility definitively, the data reported here provide little evidence that implied participants were more invested in their initial beliefs. Relative to the explicit condition, participants in the implied condition did not evidence higher levels of initial belief in the misinformation (i.e., as measured by performance in the uncorrected conditions). Second, if participants in the implied condition believed the misinformation with more conviction, one might reasonably expect that they should have been more likely to reject the correction as untrue. To the contrary, the results of Experiment 3 showed that participants in the implied condition were no more likely than those in the explicit condition to disbelieve the correction. Hence, although we cannot rule out the possibility that participants in the implied condition held their initial beliefs with more conviction, the current experiments provide no evidence in support of this account.

Another possibility is that the correction was less effective in the implied, relative to the explicit, condition because it was much more difficult for participants to successfully revise and update their initial beliefs in the implied case. Updating one's knowledge is a resource-demanding and effortful process (Ecker, Lewandowsky, Chang, et al., 2014; Ecker, Swire, & Lewandowsky, 2014; Lewandowsky et al., 2012). However, a coherent story representation that is well elaborated, highly interconnected and well structured (as in the implied misinformation case) is likely to be more difficult to revise than one where the causal links between the elements of the story are fewer and less dense. Hence, a correction following implied misinformation may have been less effective because the readers had greater difficulty updating their initial understanding.

In addition to the causal integration of the misinformation, successful correction requires detecting the inconsistency between the correction and one's initial understanding, and then using the correction to update one's knowledge. Hence, implied misinformation may have been more resistant to correction because readers may have had greater difficulty noticing the discrepancy between the correction and the misinformation when the misinformation was implied (cf. Ecker, Lewandowsky, Chang, et al., 2014). Awareness that the original misinformation has been refuted by later information is a necessary precondition for revising and updating one's original beliefs about the cause of the outcome (Marsh & Fazio, 2006; Rapp, Hinze, Kohlhepp, & Ryskin, 2014; Rapp & Kendeou, 2007, 2009; Wilson & Brekke, 1994). When reading the news story, readers construct a representation of the text, which they update while reading. If readers encounter a correction in the news story, and recognize it as such, they should attempt to revise their initial mistaken understanding and halt any further attempts to build a representation around the misinformation. To the extent that participants in the implied misinformation conditions failed to recognize that the correction was inconsistent with their initial understanding, they may not have engaged in attempts to revise their initial beliefs, and instead proceeded to build a causally coherent representation around the misleading inference, a representation that was likely augmented by their own self-generated reflections and elaborations. In contrast, participants in the explicit misinformation condition may have initiated attempts to update and undo the effects of the misinformation even while encoding the original story.

Detecting the discrepancy between the misinformation and the correction likely has other memorial consequences that may increase the effectiveness of the correction. To the extent that participants detected the discrepancy between the misinformation and the correction, they may have linked these ideas in memory, such that retrieval of the misinformation automatically cued retrieval of the correction, in a manner suggested by studies of recursive reminding (e.g., Jacoby & Wahlheim, 2013; Putnam, Wahlheim, & Jacoby, 2014; Wahlheim & Jacoby, 2013). However, to the extent that participants failed to detect the discrepancy, the misinformation and its correction will not be as strongly linked in memory (even if they coexist in memory), and recall of the misinformation may not cue retrieval of the correction. Failure to retrieve the correction when the misinformation is activated will increase misinformation's influence on performance. In sum, it is possible that as a consequence of differences in discrepancy detection, the misinformation and its correction were much more strongly linked in memory in the explicit misinformation condition than in the implied misinformation condition. By this account, the correction was less effective in the implied condition because participants more often failed to retrieve the correction when the misinformation came to mind.

As evidenced by the foregoing discussion, a variety of mechanisms may have contributed to our finding that implied misinformation was more resistant to correction than explicit misinformation. These putative mechanisms are not mutually exclusive, and may have operated in concert. It is also possible (if not likely) that, for both implied and explicit misinformation, there are individual differences in the way participants responded to the correction (see Otero & Kintsch, 1992, for evidence that subgroups of participants respond very differently to contradictions in text). Some of the mechanisms identified earlier (e.g., resistance to correction, discrepancy detection) may be associated with individual differences in personality (e.g., need for cognitive closure) and cognitive ability (e.g., executive functioning). If so, the greater persistence of implied misinformation reported here may be magnified in some populations, and minimized in others. Clearly, identifying the precise mechanisms that underlie resistance to correction following implied misinformation remains a critically important question for future research.

Conclusion

Because news reports are often released before all of the facts are known, they are sometimes mistaken. Although these mistakes may be explicit statements (e.g., an incorrect fact), in some cases the mistaken information is merely implied—such as when a news story implies that someone or something is responsible for an adverse outcome. Although participants in the implied and explicit misinformation conditions of the current experiments received nearly identical stories (that differed only in the presence/absence of a single sentence), the consistent finding was that implied misinformation was more difficult to correct than explicitly provided misinformation. Although the precise mechanisms that give rise to this finding remain to be verified by future research, the

practical implications are clear: misinformation that is "merely" implied is more difficult to eradicate than misinformation that is stated explicitly.

References

Ayers, M. S., & Reder, L. M. (1998). A theoretical review of the misinformation effect: Predictions from an activation-based memory model. *Psychonomic Bulletin & Review, 5,* 1–21. http://dx.doi.org/10.3758/BF03209454

Davies, M. F. (1997). Belief persistence after evidential discrediting: The impact of generated versus provided explanations on the likelihood of discredited outcomes. *Journal of Experimental Social Psychology, 33,* 561–578. http://dx.doi.org/10.1006/jesp.1997.1336

Duffy, S. A., Shinjo, M., & Myers, J. L. (1990). The effect of encoding task on memory for sentence pairs varying in causal relatedness. *Journal of Memory and Language, 29,* 27–42. http://dx.doi.org/10.1016/0749-596X(90)90008-N

Ecker, U. K. H., Lewandowsky, S., & Apai, J. (2011). Terrorists brought down the plane!—No, actually it was a technical fault: Processing corrections of emotive information. *The Quarterly Journal of Experimental Psychology, 64,* 283–310. http://dx.doi.org/10.1080/17470218.2010.497927

Ecker, U. K. H., Lewandowsky, S., Chang, E. P., & Pillai, R. (2014). The effects of subtle misinformation in news headlines. *Journal of Experimental Psychology: Applied, 20,* 323–335. http://dx.doi.org/10.1037/xap0000028

Ecker, U. K. H., 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,* 292–304. http://dx.doi.org/10.3758/s13421-013-0358-x

Ecker, U. K. H., 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,* 570–578. http://dx.doi.org/10.3758/s13423-011-0065-1

Ecker, U. K. H., Lewandowsky, S., & Tang, D. T. W. (2010). Explicit warnings reduce but do not eliminate the continued influence of misinformation. *Memory & Cognition, 38,* 1087–1100. http://dx.doi.org/10.3758/MC.38.8.1087

Ecker, U. K. H., Swire, B., & Lewandowsky, S. (2014). Correcting misinformation: A challenge to education and cognitive science. In D. N. Rapp & J. L. G. Braasch (Eds.), *Processing inaccurate information: Applied and theoretical perspectives from cognitive science and the educational sciences* (pp. 13–38). Cambridge, MA: MIT Press.

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods, 39,* 175–191. http://dx.doi.org/10.3758/BF03193146

Gallo, D. A. (2010). False memories and fantastic beliefs: 15 years of the DRM illusion. *Memory & Cognition, 38,* 833–848. http://dx.doi.org/10.3758/MC.38.7.833

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and Semantics: Vol. 3. Speech Acts* (pp. 41–58). New York, NY: Academic Press.

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,* 73–81. http://dx.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 & Cognition, 2,* 201–209. http://dx.doi.org/10.1016/j.jarmac.2013.10.001

Hart, P. S., & Nisbet, E. C. (2012). Boomerang effects in science communication: How motivated reasoning and identity cues amplify opinion polarization about climate mitigation policies. *Communication Research, 39,* 701–723. http://dx.doi.org/10.1177/0093650211416646

Hilton, D. J., McClure, J., & Sutton, R. M. (2010). Selecting explanations from causal chains: Do statistical principles explain preferences for voluntary causes? *European Journal of Social Psychology, 40,* 383–400.

Jacoby, L. L., & Wahlheim, C. N. (2013). On the importance of looking back: The role of recursive remindings in recency judgments and cued recall. *Memory & Cognition, 41,* 625–637. http://dx.doi.org/10.3758/s13421-013-0298-5

Johnson, H. M., & Seifert, C. M. (1994). Sources of the continued influence effect: When misinformation in memory affects later inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 1420–1436. http://dx.doi.org/10.1037/0278-7393.20.6.1420

Johnson, H. M., & Seifert, C. M. (1999). Modifying mental representations: Comprehending corrections. In H. van Oostendorp & S. R. Goldman (Eds.), *The construction of mental representations during reading* (pp. 303–318). Mahwah, NJ: Erlbaum.

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin, 114,* 3–28. http://dx.doi.org/10.1037/0033-2909.114.1.3

Kim, S. I., & Van Dusen, L. M. (1998). The role of prior knowledge and elaboration in text comprehension and memory: A comparison of self-generated elaboration and text-provided elaboration. *The American Journal of Psychology, 111,* 353–378. http://dx.doi.org/10.2307/1423446

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,* 106–131. http://dx.doi.org/10.1177/1529100612451018

Lewandowsky, S., Stritzke, W. G. K., Oberauer, K., & Morales, M. (2005). Memory for fact, fiction, and misinformation: The Iraq War 2003. *Psychological Science, 16,* 190–195. http://dx.doi.org/10.1111/j.0956-7976.2005.00802.x

Marsh, E. J., & Fazio, L. K. (2006). Learning errors from fiction: Difficulties in reducing reliance on fictional stories. *Memory & Cognition, 34,* 1140–1149. http://dx.doi.org/10.3758/BF03193260

Mason, R. A., & Just, M. A. (2004). How the brain processes causal inferences in text. *Psychological Science, 15,* 1–7. http://dx.doi.org/10.1111/j.0963-7214.2004.01501001.x

McNamara, D. S., & Magliano, J. P. (2009). Toward a comprehensive model of comprehension. In B. Ross (Ed.), *The psychology of learning and motivation* (Vol. 51, pp. 297–384). New York, NY: Elsevier Science. http://dx.doi.org/10.1016/S0079-7421(09)51009-2

Moscovitch, M., & Melo, B. (1997). Strategic retrieval and the frontal lobes: Evidence from confabulation and amnesia. *Neuropsychologia, 35,* 1017–1034. http://dx.doi.org/10.1016/S0028-3932(97)00028-6

Myers, J. L., Shinjo, M., & Duffy, S. A. (1987). Degree of causal relatedness and memory. *Journal of Memory and Language, 26,* 453–465. http://dx.doi.org/10.1016/0749-596X(87)90101-X

Nyhan, B., & Reifler, J. (2010). When corrections fail: The persistence of political misperceptions. *Political Behavior, 32,* 303–330. http://dx.doi.org/10.1007/s11109-010-9112-2

Nyhan, B., Reifler, J., & Ubel, P. A. (2013). The hazards of correcting myths about health care reform. *Medical Care, 51,* 127–132. http://dx.doi.org/10.1097/MLR.0b013e318279486b

Otero, J., & Kintsch, W. (1992). Failures to detect contradictions in a text: What readers believe versus what they read. *Psychological Science, 3,* 229–235. http://dx.doi.org/10.1111/j.1467-9280.1992.tb00034.x

Putnam, A. L., Wahlheim, C. N., & Jacoby, L. L. (2014). Memory for flip-flopping: Detection and recollection of political contradictions. *Memory & Cognition, 42,* 1198–1210. http://dx.doi.org/10.3758/s13421-014-0419-9

Rapp, D. N., Hinze, S. R., Kohlhepp, K., & Ryskin, R. A. (2014). Reducing reliance on inaccurate information. *Memory & Cognition, 42,* 11–26. http://dx.doi.org/10.3758/s13421-013-0339-0

Rapp, D. N., & Kendeou, P. (2007). Revising what readers know: Updating text representations during narrative comprehension. *Memory & Cognition, 35,* 2019–2032. http://dx.doi.org/10.3758/BF03192934

Rapp, D. N., & Kendeou, P. (2009). Noticing and revising discrepancies as texts unfold. *Discourse Processes, 46,* 1–24. http://dx.doi.org/10.1080/01638530802629141

Schacter, D. L., Norman, K. A., & Koutstaal, W. (1998). The cognitive neuroscience of constructive memory. *Annual Review of Psychology, 49,* 289–318. http://dx.doi.org/10.1146/annurev.psych.49.1.289

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

Wahlheim, C. N., & Jacoby, L. L. (2013). Remembering change: The critical role of recursive remindings in proactive effects of memory. *Memory & Cognition, 41,* 1–15. http://dx.doi.org/10.3758/s13421-012-0246-9

Wegner, D. M., Wenzlaff, R., Kerker, R. M., & Beattie, A. E. (1981). Incrimination through innuendo: Can media questions become public answers? *Journal of Personality and Social Psychology, 40,* 822–832. http://dx.doi.org/10.1037/0022-3514.40.5.822

Wilkes, A. L., & Leatherbarrow, M. (1988). Editing episodic memory following the identification of error. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology, 40,* 361–387. http://dx.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 A: Human Experimental Psychology, 52,* 165–183. http://dx.doi.org/10.1080/713755808

Wilson, T. D., & Brekke, N. (1994). Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin, 116,* 117–142. http://dx.doi.org/10.1037/0033-2909.116.1.117

Appendix A

Experiment 1 News Story

Message 1	At 3:00 p.m., May 2nd, police responded to a call made from a home on Acorn Street in the middle-class, residential neighborhood of Sunny Hollow.
Message 2	The homeowner, June Harter, reported that her jewelry box was missing. The box's contents included gold chains, gold and silver earrings, rings, and pendants with precious stones.
Message 3	She discovered that the box was missing when she and her husband returned from a vacation and she wanted to put a new necklace she had bought in the jewelry box. The box had been stored in a locked drawer in her bedroom dresser.
Message 4	*All explicit and implied conditions:* Mrs. Harter swears that she had checked the box before leaving on vacation and that everything was in order. A tall tree arches near the bedroom window, but police have found no evidence of tampering with the window. *Control:* Mrs. Harter specifically remembers that she went through her jewelry before leaving to get her pearl earrings. They were the last things she packed before the Harters' son, Evan, arrived to take them to the airport.
Message 5	*Explicit:* Police suspect that the Harters' son, Evan, may have taken the box from the house. The Harters report that they had asked their son to check in on the house periodically during their absence. The son also did other odd jobs for many of the neighbors to pay off his recent gambling debts. *Implied:* The Harters report that they had asked their son, Evan, to check in on the house periodically during their absence. The son also did other odd jobs for many of the neighbors to pay off his recent gambling debts. *Control:* A tall tree arches near the bedroom window, but police have found no evidence of tampering with the window.
Message 6	Sunny Hollow has been hit with a number of thefts recently, but there are no arrests or leads in these cases so far.
Message 7	The Harters' next-door neighbor reported that she noticed a light on in the Harters' house after her dog suddenly began barking late Saturday evening, April 28th. An unfamiliar dark-colored car had been parked in a nearby alley.
Message 8	However, a search for footprints and tire tracks has turned up inconclusive due to a recent rainstorm. In the course of the investigation, an officer also noted a broken latch on a basement window.
Message 9	Police are still attempting to determine whether other valuables are missing from the home. The television and home computer, however, had not been disturbed.
Message 10	The Harters reported that they had contacted their insurance company about the loss. The last appraisal showed the box's content to be worth several thousand dollars.
Message 11	*Corrected explicit and corrected implied conditions:* Police have now confirmed that Evan had been called away on business and had not been in town to look after the house during the Harters' vacation. *Corrected control condition:* Police have now confirmed that Evan had been called away on business and had not been in town during the Harters' vacation.
Message 12	In addition, Mrs. Harter is considering offering a reward for the return of several of the pieces of jewelry because they have great sentimental value for her. She says that there would be no questions asked.
Message 13	Detectives will look for similarities between this case and the other thefts reported in the neighborhood recently. If you have any information which may aid in the investigation, please contact the police department.

(*Appendices continue*)

Appendix B

Final Questionnaire

Causal Question

1. What caused the box to be missing from the Harters' home?

Factual Questions

2. How much did an appraisal show the box's contents to be worth?
3. Where was the Harters' home located?
4. Where was the jewelry box normally kept?
5. Why did Mrs. Harter consider offering a reward?
6. What did the Harters' next-door Neighbor notice?
7. What kinds of jewelry did the box contain?
8. When did Mrs. Harter discover that the jewelry box was missing?
9. What did the police notice about the bedroom window?
10. When did the neighbor's dog suddenly start barking?

Inference Questions

11. Why might the neighbor's dog have been barking?
12. Whose car might the neighbor have noticed parked in the alley?
13. Why might the son feel bad about the incident?
14. What could the Harters have done to better avoid this problem?
15. What steps should the police take next?
16. Why wasn't the television taken?
17. How might the thief have gotten into the house?
18. How do you think the thief got into the locked drawer to steal the jewelry box?
19. What might be responsible for the other thefts in the neighborhood recently?
20. Who, if anyone, should be questioned more thoroughly by the police?

Direct Measure of Belief in the Son's Involvement

21. How much do you believe that Evan Harter was involved in the theft of the jewelry? Please indicate your answer using the scale below:

1	2	3	4	5	6
Highly unlikely	Moderately unlikely	Somewhat unlikely	Somewhat likely	Moderately likely	Highly likely

Free Recall Question

22. Now, use the space below to describe the event reported in the news reports you read. Write down everything you remember *in as much detail as possible* but also be *as accurate as possible*.

Manipulation Check Questions

23. (Corrected Conditions Only) What did the story report about where Evan Harter was during the Harters' vacation?

24. What facts about the case did the police change their minds about, based on information they discovered later?

Appendix C

Experiment 2 News Story

Messages 1–10	Messages identical to Experiment 1.
Message 11A	Correction messages identical to Experiment 1: *Corrected explicit and corrected implied conditions:* Police have found that Evan had been called away on business and had not been in town to look after the house during the Harters' vacation. *Corrected control condition:* Police have found that Evan had been called away on business and had not been in town during the Harters' vacation.
Message 11B	*All corrected conditions:* They have identified a suspect, ex-convict Dan Fowler. Fowler recently sold expensive jewelry to a pawn shop, and his girlfriend, Sarah, works for a cleaning service that regularly cleans the Harter's home.
Messages 12–13	Messages identical to Experiment 1.

Received June 6, 2014
Revision received May 15, 2015
Accepted May 16, 2015 ■