

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.

The Continued Influence of Misinformation in Memory:
What Makes a Correction Effective?

Colleen M. Seifert
University of Michigan
525 East University
Ann Arbor, MI 48109-1109
(734) 763-0210; seifert@umich.edu

- I. The Continued Influence of Misinformation in Memory
- II. Sources of the Continued Influence Effect
 - A. Failing to edit inferences in memory
 - B. Processing contradiction
 - C. Accessing misinformation in memory
 - D. Replacing misinformation
 - E. Covering the “causal gap”
- III. The Pragmatics of Correcting Misinformation
 - A. Conversational conventions
 - B. Explanatory corrections
 - C. Qualifying inferences
- IV. Effective Corrections of Misinformation

Acknowledgements: Many thanks to my collaborators on the studies described in this chapter, Julie G. Bush and Hollyn M. Johnson. Thanks also to several research assistants helped us conduct the studies, including Christy Marshuetz, Jennifer Coe, Andrew Jeung and Srinika Narayan. Brian Ross provided helpful comments on an earlier draft.

The Continued Influence of Misinformation in Memory

In a dynamic world, information in memory is frequently outdated, corrected, or replaced. People often make use of this misinformation in memory during later reasoning (Loftus, 1978). Some studies have shown that reasoners can draw upon information in memory even though they also acknowledge it as incorrect (Wilkes & Leatherbarrow, 1988; Johnson & Seifert, 1994). What makes information in memory so compelling that it outlives its validity? Consider this series of news reports:¹

Monday: “A family of four was found dead in their home. The family had dined the night before at the "China Garden" restaurant.”

Tuesday: “The medical examiner now reports that food poisoning has been ruled out as a possible cause of death for the family.”

Saturday: “"China Garden" has gone out of business.”

This example demonstrates several important features of this phenomenon. First, the information and its correction are clearly presented and connected together in memory. Second, the information and its correction are believable, and accepted as valid in the absence of any conflicting information. Third, following correction, the value of the initial information (as incorrect) is clearly identified. And finally, the misinformation and its correction are both accurately recalled later, at the time of reasoning. Yet despite these factors, the misinformation continues to influence later judgments after correction, a phenomenon called the “continued influence effect” (Johnson & Seifert, 1994).

Most prior research on misinformation in memory has focused on the effect identified by Loftus (1978). In this paradigm, conflicting information is introduced following the initial acquisition of an account. For example, after watching a video of a traffic accident where a “stop” sign is depicted, subjects are asked how fast the car was traveling when it passed the “yield” sign. Using post-event questions, new misinformation is introduced into memory for the original account. As a result, people sometimes reported the wrong information (the “yield” sign) when later asked about the account. In these misinformation effects, people appear to be unaware of the conflict between the initial account and the later misleading question (Loftus, 1979). When asked about the “yield” sign, they don’t protest that it was in fact a stop sign; instead, they answer the question, and then sometimes incorporate the “yield” sign into their account of the original incident. Later work attempted to identify the effect as retrieval-based (McCloskey & Zaragoza, 1985), where two conflicting representations remain in memory and can be accessed separately.

While the content of the misinformation may play an important role, prior studies typically changed a more minor detail in a scene, such as the brand of a can of soda (Belli, 1989). Other research has demonstrated that readers can be aware of an inconsistency (e.g., between an earlier described characteristic of a protagonist and a subsequent target action) (Albrecht & O’Brien, 1993). Despite awareness, readers may not integrate updated information with the initial account in memory, but may wait to incorporate the changes only when needed in reasoning (McKoon &

Ratcliff, 1992; O'Brien, Rizzella, Albrecht, & Halleran, 1998). These studies suggest that whether or not the reader is aware of contradiction, they may add updated information into memory without reference to the initial representation.

However, this “multiple representation” account cannot address the example outlined above. In the restaurant account, the misinformation and its correction are clearly connected together in memory. People accurately report that food poisoning was not the cause of death; yet, they continue to feel wary of eating in the same restaurant. To explain this phenomenon, we will have to account for the continued influence of misinformation in the presence of direct correction. This phenomenon has been demonstrated in a variety of tasks, including inferences drawn from text (Wilkes, 1988; Wilkes & Leatherbarrow, 1988; Wilkes & Reynolds, 1999), forming impressions of individuals (Golding, Fowler, Long, & Latta, 1990; Fleming & Arrowood, 1979; Wyer & Budesheim, 1987; Wyer & Unverzagt, 1985), and creating theories from data (Ross, Lepper, & Hubbard, 1975; Anderson, Lepper, & Ross, 1980; Anderson, 1982; Anderson, 1983; Anderson, New, & Speer, 1985).

II. Sources of the Continued Influence Effect

How is it possible that people reading a direct correction continue to use misinformation? As measured by their answers to reasoning questions, corrected subjects did not differ from those who never received a correction. This suggests the factors driving the continued influence effect must be powerful, such that subjects can find no alternative to using information that is incorrect. In the next sections, we consider several alternative explanations for the continued influence effect, including a failure to edit inferences in memory, the availability of the misinformation in memory, and the role of a “causal gap” in information. What factors influence reasoners to use information they know to be inaccurate?

A. *Failing to edit inferences in memory*

One account suggests the later use of misinformation involves inferences drawn from the information before it could be corrected. In this “editing” hypothesis, one might use information to draw a conclusion about an event; later, when the information is corrected, the conclusion based on it is not accessed or corrected. So, elaborations based on misinformation may continue unaltered in memory, and influence later judgments (Loftus, 1979; Wilkes & Leatherbarrow, 1988; Ross, Lepper, & Hubbard, 1975).

Wilkes and Leatherbarrow (1988) presented a clever experimental paradigm that demonstrated the continued influence of corrected information in memory. Subjects were presented a series of “on the scene” news reports regarding a warehouse fire in progress (Millis & Erdman, 1998). The fifth message in the series noted that “cans of oil paint and gas cylinders” had been stored in a room where the fire started. Later messages referred to “oily smoke and sheets of flame,” “explosions,” and “toxic fumes” as features of the fire. Near the end of the account, a message stated a correction:

that no paint or gas cylinders had been in the storage room. After a ten minute distractor task, subjects were asked factual (What business was the firm in?) and “inference” (What was a possible cause of toxic fumes?) questions about the account. If readers grasped the correction, they should omit any mention of the paint and gas when questioned about possible causes of the fire. These results showed that subjects who saw the corrected information were much more likely to refer to it than were control subjects who never saw the initial misinformation. In fact, those who saw the correction did not differ statistically in the number of responses involving the discredited materials from those who saw the information without any correction!

Wilkes and Leatherbarrow (1988) suggest that people failed to edit elaborative inferences made during reading before the correction occurred. As subjects read further details about the fire, they may have incorporated the volatile materials in the storage closet as explanations for the fire’s observed features. For example, when reading about the explosions that occurred, they may have considered the “pressurized gas cylinders” as a likely source. Later in the story, the correction appears and clearly rules out the presence of gas cylinders in the storage closet. While able to correct the misinformation, subjects may be unable to find all prior inferences based on it, and edit them from memory. Consequently, when asked what might have caused the explosions, they may fall back on the “gas cylinders” explanation generated before the correction occurred. Similarly, studies such as Ross, Lepper, and Hubbard (1975), found that stereotypes of an individual based on behavior descriptions were not corrected when the descriptions were. Due to failures in editing, theories appear to outlast the discrediting of the data they are based upon.

An alternative account suggests that misinformation could be used in later inferences even after it had been corrected. Much later after reading, perhaps even at the time of test, people may refer back to information they know is inaccurate, and yet use it to draw new inferences. This would appear irrational; however, some reasoning situations may lead people to make use of information they already know to be incorrect. For example, if no other alternatives can be generated, people may fall back on past accounts even though they are known to be inaccurate.

To test the “failure to edit” hypothesis, Hollyn Johnson and I conducted a study designed to vary the opportunities for drawing further inferences before correction (Johnson & Seifert, 1994). In the Wilkes and Leatherbarrow (1988) paradigm, subjects read the bulk of the story before receiving the correction. In our study, we compared this condition to one where the correction appeared immediately after the misinformation. That is, subjects read about the storage room contents containing volatile materials, and then immediately read that the room had actually been empty. As these subjects went on to read about the fire, they would be presumably less able to draw on the potential explanation of the volatile materials in the storage room. Since these subjects already know the information is incorrect, they may perform more like subjects who never saw the misinformation, and be forced to create other explanations for the fire’s features. We know it is possible to answer the questions without referring to the volatile materials, since some subjects in the “no reference” control group never saw the initial information about them.

We manipulated where the correction was presented (message number 6 in the early correction group, and message number 11 (of 13) in the late correction group); otherwise, the materials (see Appendix A), procedure, and scoring were as in Wilkes and Leatherbarrow (1988). Open-ended responses to fact and inference questions (see Appendix B) were scored on two measures. The “volatile materials” score included the number of responses consistent with believing that the volatile materials caused or contributed to the fire. Examples of positive responses for this scale include:

11. *Why did the fire spread so quickly? “Oil fires are hard to put out.”*
13. *What was the possible cause of the toxic fumes? “Burning paint.”*
17. *Why do you think the fire was particularly intense? “The pressurized cylinders.”*

A second measure counted the number of direct and uncontroverted references to the volatile materials themselves (paint cans or gas cylinders) in any of the responses. Presumably, if the correction were incorporated into the account in memory, responses would involve few direct mentions of the volatile materials.

Surprisingly, our results showed that providing an immediate correction (in the message following the target information) resulted in no improvement in the use of misinformation compared to subjects who saw the correction at the end of the story. Both groups referred to the volatile materials (paint cans and gas cylinders) as potential causes of fire features (such as explosions and toxic fumes). Since the immediate correction group did not have the opportunity to draw further inferences before the correction occurred, these results suggest people drew these inferences using the misinformation after the correction had occurred. While failures to edit inferences may account for misinformation effect in other circumstances, we found that people may also draw new inferences explicitly using information already known to be incorrect. This surprising finding was replicated by Wilkes and Reynolds (1999), and suggests that inferences during reading play a much smaller role in the continued influence effect than inferences made on demand at question time. The continued influence effect may be a circumstance where initial processing involves a “minimalist” approach (McKoon & Ratcliff, 1992), where effortful inferences are postponed unless needed for local coherence. It appears that readers waited until pressed by questions to attempt to resolve the contradiction of the initial information and the later correction.

B. Processing contradiction

The fact that the time of correction (immediate vs. delayed) had no effect on these responses suggests that the continued influence of misinformation is occurring well after correction. Subjects appear to be acknowledging the correction, yet at the same time, using the information that was corrected. The contradictory nature of this reasoning is demonstrated by subjects’ protocols: on a single response sheet, one subject wrote the following:

13. *What was the possible cause of the toxic fumes?*
“The paint and gas cylinders.”
21. *What was the point of the second message from the police?*

“That the storage room didn’t have any gas or paint, but was empty.”

Subjects in the correction group often gave contradictory responses when asked about the facts of the fire vs. the cause of the fire. For example, consider this protocol:

12. *For what reason might an insurance claim be refused?*

“Because of the oil paint and the pressurized gas cylinders that were in the closet.”

21. *What was the point of the second message from the police?*

“To state that the gas cylinders and the oil paint cans were not in the closet.”

These contradictory responses within subjects’ protocols did not arise from only a few subjects. Across the board, 90% of subjects who saw the correction made references to “gas cylinders” and “paint cans” in their responses, while only 20% of control subjects (who never saw the initial information) did so. This suggests that almost all who heard the information were influenced by it despite the correction.

How might subjects reconcile their belief that the correction was accurate (there were no volatile materials in the storeroom) with their belief that important features of the fire (toxic fumes, explosions, intense heat) must have been caused by these same materials? Subjects may have taken a variety of turns in interpreting the corrections (Baker & Wagner, 1979; Schul & Bernstein, 1985). First, did subjects understand the correction? In an attempt to clarify the correction message, we ran a replication where the correction was stated more directly:

Correction message (from Wilkes & Leatherbarrow, 1988):

Message 11: 10:30 a.m. A second message received from Police Investigator Lucas regarding the investigation into the fire. It stated that the closet reportedly containing cans of paint and pressurized gas cylinders had been actually been empty before the fire.”

Alternative correction message (from Johnson & Seifert, 1994):

Message 11: 10:30 a.m. A second message received from Police Investigator Lucas regarding the investigation into the fire. It stated that there were no cans of paint or gas cylinders in the closet that had reportedly contained them; the closet had actually been empty before the fire.

The results were identical for the two correction messages (Johnson & Seifert, 1994). This suggests subjects are interpreting the correction accurately, as well as grasping the intent of the correction. Almost all subjects correctly reported that the storage closet did not contain the paint cans and gas cylinders as initially reported. Around 90% accurately recalled the correction, and omitting those who did not from the analyses did not change the results. Subjects also appeared to understand the story itself equally well in control and correction accounts, in that the groups did not differ in memory for story facts, or in their ability to summarize the reports.

Second, did subjects believe the correction information? The message stated that the room was empty; however, subjects could have decided that this message was in error or intentionally misleading, and held to the belief that the materials were in fact in the storage room (Baker & Wagner, 1987). If they did so, they would have the volatile materials explanation to use in

answering the inference questions. However, none of the subjects gave responses consistent with this scenario. Since the responses were free form, subjects could have modified or “hedged” their responses (the closet was “supposedly” empty; or, “the police thought the closet was empty.”) Instead, subjects answered questions about the storage room contents in a manner consistent with the belief that it was indeed empty. And less than 20% of subjects proposed that the volatile materials were responsible for the fire when asked directly. But when questions required a causal explanation, subjects fell back on the misinformation. The volatile materials were mentioned on inference questions about the fire, such as “Why do you think the fire was particularly intense?”

Perhaps subjects simply could not reconcile the two views, and kept both in memory (Hertel, 1982). This protocol suggests the subject did not choose which account to believe, but simply kept track of what was reported:

13. Were you aware of any contradictions in the reports that you read?

“Yes, because at first the cylinders and cans were in the closet and then they weren’t.”

If unable to make sense of the need for the explanation provided by the volatile materials and the presence of the correction, subjects might choose to use the information despite the contradiction. Or, perhaps they generated alternative versions of the story that incorporated both the information and the correction. They might attempt to alter their accounts to posit the presence of the volatile materials elsewhere in the warehouse, where they could have contributed to the fire. For example, they could interpret the correction to mean that the materials were not in the storeroom; however, they might be in another nearby room in the building. Such a combination of ideas still represents the influence of misinformation because this interpretation is not suggested by the direct correction. Subjects may have drawn other inferences in order to preserve both the accuracy of the correction and the causal power of the misinformation.

In other studies, we examined the process of correction during reading using a speeded verification task (Johnson & Seifert, 1998). In these studies, subjects read short news stories containing corrections of fact. For example, a hockey player from Philadelphia was reported to play for a team in Boston; later, this was corrected by stating he had actually played for New York. The last line of the story ended with a reference to “the team.” Immediately, a test phrase was presented, and subjects had to verify whether the item appeared in the story they just read. The results were surprising: Subjects were faster to verify the correct referent (New York) than a filler (Philadelphia), but were just as fast to verify the misinformation (Boston). In fact, both referents (the initial and the corrected) showed facilitation even after a delay. These findings suggest that the memory representation following correction includes an association with the invalidated information even when subjects correctly report only the valid referent when questioned.

In sum, subjects processing a correction face a contradiction: they know the information is not correct, but they also know the information. In responding to questions, their inferences are simply not the same as subjects who were never exposed to the misinformation. Control subjects who never saw a message about volatile materials were also required to answer the inference questions about possible causes of toxic fumes and explosions. But the correction subjects are different in that they

had seen a good explanation for those features in the misinformation presented. Perhaps the mere presence of the information in the account, even though later corrected, is enough to lead to the use of the information in later reasoning.

C. Accessing misinformation in memory

One account of information use argues that the mere presence of information in memory may lead to its incorporation in later reasoning. If information is available in memory, regardless of knowledge about its status, it may affect the conclusions drawn (Tversky & Kahneman, 1973). This “accessibility” argument (Wegner & Erber, 1992) emphasizes the context of currently active information as a major factor in how information is used, and posits no role for knowledge about the status of information. In this account, the mere mentioning of information increases its availability and subsequent use in reasoning. To examine this possibility, we compared corrected and uncorrected story versions with a version where the volatile materials (paint cans and gas cylinders) were mentioned only incidentally in the story (the fire was reported by another store selling paint and gas). We also included a group of subjects who never saw the information about the volatile materials, but who completed an unrelated intervening word generation task where “gas” and “paint” were presented.

Finally, we considered whether the availability of alternative causes would help subjects to avoid the use of the misinformation. We asked another group of subjects to generate their own potential causes for the fire after completing the account (Johnson & Seifert, 1999). These subjects received the information and the correction, and then were asked to generate “4 or 5 possible causes for a fire of the type they had just read about.” Perhaps the availability of alternative causes would help to avoid the influence of the corrected information. If the mere availability of information in memory results in its use, introducing alternatives they generated themselves may help subjects to create other inferences about the fire’s cause in the warehouse story.

The results showed that the mere availability of the information was not sufficient to promote its use in reasoning. As in the earlier studies, the results showed the continued influence effect: those who saw a correction (stating that the storeroom had actually been empty) were just as likely to report the volatile materials as potential causes as subjects who saw the information but never received a correction. However, subjects were likely to refer to those materials as explanations for the fire’s features only when they saw the volatile materials presented in the story as a possible cause of the fire. Mere mention of volatile materials in an incidental role or during the experiment had no effect on responses. Some subjects who saw the initial information and correction were also asked to generate other causes for the fire, and they suggested arson, electrical malfunctions, lightning, bombs, cigarettes, gas leaks, and many others. However, generating these alternatives did not decrease the likelihood of mentioning the volatile materials to answer inference questions. Subjects only incorporated the misinformation in their reasoning when it was suggested initially as a potential cause. However, generating other potential causes for the fire did not alleviate subjects’ use of the misinformation. Together, these findings suggest the source of the continued influence of

misinformation may lie in its causal role within an account. That is, the effect is not due to just any misinformation, nor to the absence of a cause: Perhaps the power of the misinformation in these studies lies in its causal role within the account.

D. Replacing misinformation

Consider the subject who never saw the “volatile materials” information. This subject must answer questions about “explosions” and “toxic fumes” in the absence of any clear cause for those fire features. The other messages about the fire contain more details, but suggest no other source of materials that would lead to such dramatic consequences. For these subjects, the questions about the cause of the fire lack any compelling explanation from the story itself. These subjects give responses typically involving other presumed contents of a “stationery warehouse;” for example, some subjects said that the toxic fumes may have been caused by “toner fluid for copiers,” or that the intense heat of the fire may arise from “large amounts of paper stored” in the warehouse. While these responses are adequate, they lack a compelling causal connection to the dramatic features (e.g., explosions) observed. Causal connections are obviously important in understanding a series of events like the warehouse fire, and causality is an indicator of the importance of an event within an account (Trabasso, & Sperry, 1985). The correction group could avoid mentioning the volatile materials, and instead provide answers similar to the “no mention” group. However, having heard the misinformation, correction subjects have seen a better causal explanation that “covers” the features observed. Their only problem is that they know the explanation is not correct.

There is some suggestion that causal connections between the discredited information and other story elements are central to the continued influence of misinformation. Wilkes (1988) showed that discredited information that played a central role in the account, such as the origin of the fire in the warehouse story, played a bigger role in later inferences. In contrast, when the discredited information did not play a central role, it appeared less often in later responses. Wilkes and Reynolds (1999) varied the number of fire features related to the misinformation, and found that even just one unexplained causal feature resulted in illegitimate inferences using the discredited information. Their studies suggest that the number of causal links related to the negated information is not a key factor in misinformation use; rather, the need for a compelling explanation for a single feature will result in continued influence of misinformation. In the fire story, the presence of the volatile materials covers a “gap” between what is known about the warehouse and the fire’s features that must be explained. The features suggest some materials on the scene were capable of generating great heat, toxic fumes, and explosions. If no other source for these features is presented, the misinformation about volatile materials in the storeroom remains the best explanation for the features observed. How does the individual reconcile accuracy with the need for causal completeness? In the absence of alternatives, they may choose to ignore the inaccuracy, and consider the causal story as “too good to not be true.” An accurate story without a cause leaves a gap in understanding, and the gap may be more bothersome when a good explanation has been presented and ruled out.

Perhaps people fall back on misinformation because there is no alternative that provides the same causal coverage. One way to test this is to provide a replacement cause; then, subjects who were troubled by the lack of a potent cause will have an alternative to adopt, and may then avoid falling back on the misinformation. In another study, we followed the correction with a later message suggesting an alternative cause for the fire. This message suggested that arson might have been involved due to “evidence of gasoline-soaked rags near where the bales of paper had been stored in the storage hall, as well as several emptied steel drums of suspicious nature.” Misinformed subjects could then avoid use of the “volatile materials” information by drawing on this alternative cause for the source of explosions and toxic fumes. On the other hand, the misinformation may be playing such a powerful role because of its primacy in the account. The presence of the misinformation, occurring early in the story, may result in a representation from which the initial information cannot be discarded. If so, we would expect to see references to the volatile materials even in the presence of a later alternative.

A control group saw the arson information only, and not the original “volatile materials” information. Their results showed that the arson explanation alone works as well as the volatile materials in answering the inference questions. We also replicated the continued influence effect, where the “information only” group and the “information plus correction” group again made equally frequent mentions of the volatile materials. However, when the information was corrected and then later replaced by the arson information, subjects were now less likely to fall back on the information to explain the fire. Instead, they called upon the arson materials to account for the features of the fire. This was the first comparison where a group that saw a correction performed as well as those who never saw the information. Successful correction of misinformation finally occurred, but only when a causal alternative to the misinformation was provided later in the story. Correction subjects still made references to the volatile materials in their answers; however, they were less likely to refer to them as potential causes, and more likely to refer to the arson materials as the source of the fire’s features.

E. Covering the “causal gap”

Successful correction of misinformation appears to require an alternative causal account. The alternative appears to replace the misinformation by providing a potential cause for the events that is equally satisfactory in explaining the circumstances. But what is it about the alternative provided that led to its acceptance by subjects? Would any alternative, no matter how relevant, be adopted? Or does the success of correction depend on the replacement’s ability to account for the same features?

In a sense, we are asking, “What makes a cause compelling?” What type of alternative will be accepted as a replacement for the role of the volatile materials in the fire? In the generation experiment above, asking subjects to think of possible alternatives did not decrease their likelihood of falling back to the volatile materials. One possibility is that a particular cause generated may not be adequate to explain the fire’s features. Subjects generated a wide variety of possible causes

varied (e.g., cigarette ashes, lightning strikes, even spontaneous combustion); however, if none of them could account for the fire's features, the volatile materials misinformation may remain a better alternative.

Perhaps the efficacy of the arson replacement is due to its ability to fill a causal role. The arson information provided included similar materials that could account for toxic fumes and explosions. Because it appears to "cover" the questions asked, this replacement may be effective in stemming the use of the misinformation. This suggests that a replacement cause that fails to account for the fire's features would be less successful. For example, what if the replacement message instead suggested "evidence of water-soaked rags near where the bales of paper had been stored in the storage hall, as well as several steel drums of water?" This alternative seems unlikely to explain the features of the fire; as a consequence, subjects may again be forced to fall back on the misinformation.

At the least, if a cause is mentioned in the text as being present on the scene, it may have greater validity. For example, in the "incidental mention" condition, the volatile materials were in a neighboring store and not in the warehouse; consequently, the materials were not seen as possible causes for the warehouse fire. Perhaps the replacement cause must be instantiated on the scene, as was the initial misinformation. While subjects in the generation experiment came up with many varied causes for fires, they failed to adopt any of the scenarios they generated as viable alternatives to the misinformation. Perhaps they were unwilling to introduce explanations that are not connected to the specific information provided in the account. The misinformation may have to be replaced by an explanation with equally strong causal coverage.

Another aspect of the arson replacement that may facilitate its adoption is its plausibility. If the alternative cause offered is unlikely to occur, it may be less acceptable to readers, and may result in continued influence from the misinformation. For example, if the replacement cause is "electrical short circuit," one may adopt that alternative more readily than "spontaneous combustion." In the generation study, subjects named typical fire sources (cigarette discarded, wiring, fireworks) that varied in their plausibility as the cause of a warehouse fire.

In another experiment, we compared the nature of the replacement cause provided (Johnson & Seifert, 1999). All subjects saw the misinformation, the correction, and later, the arson alternative. However, we varied the materials left in the storage room by the arsonists. We selected materials that had either a high or a low likelihood of being present at the warehouse; in a sense, this reflected the plausibility of the cause. The presence of gasoline-soaked rags and boxes of pencils fits with the story of arson in a stationery warehouse. The presence of cartons of hair spray and bolts of cloth seem less plausible, since they don't appear connected to either an arson attempt or to the typical contents of a stationery warehouse. In addition, we selected materials that offered either high or low coverage of the causal features in the warehouse fire. These were designed to provide either causal coverage or inadequate explanation of the fire's events. For example, the gasoline-soaked rags and cases of hairspray provide explanations for some of the fire features described later, such as "toxic fumes" and "explosions." Other materials, such as "school supplies" and "bolts of cloth," fail to

provide any connection to the fire's features. This design allows us to compare alternatives on coverage and plausibility, and identify the factors involved when a replacement cause is successfully adopted.

All four groups in the study saw the fire story with the initial mention of the volatile materials in Message 5, a correction in Message 11, and a replacement (other materials in the warehouse storeroom) in Message 12. The likelihood manipulation varied how plausible the item was for a stationery warehouse, and the coverage manipulation varied how well the alternative explained the features of the fire. The results showed no main effect of plausibility; that is, whether the materials were typical or atypical for the warehouse setting did not affect whether subjects used the misinformation in their responses. However, the conditions where the replacement provided high causal coverage did decrease the use of the misinformation. If the replacement cause provided an explanation for the fire's features, that cause was used in the responses to questions, resulting in less frequent use of the misinformation.

In addition, an interaction occurred between the two factors, suggesting that in high coverage conditions, plausibility played little role. However, if the alternative provided offered low coverage of the fire features, the more plausible alternatives were adopted more often. In the absence of a compelling alternative, a likely one may work better to avoid falling back on misinformation. But the most effective replacement for misinformation appears to require coverage of the causal features. With a powerful causal explanation, even if implausible or unlikely, participants were less likely to fall back on the misinformation when answering questions about the account.

In sum, these studies show that people sometimes continue to fall back on information they know is not correct in order to make inferences and explain events. The continued influence effect appears to involve using information during reasoning that has already been discredited, and is motivated by the need for causal explanations for events in the story. When required to make inferences about the causes of events, people may fall back on explanations that have been discredited but that fill a causal gap in the account. By providing a replacement cause with the correction, the influence of the discredited information can be avoided. This suggests correction can be successful, and people can avoid using initial misinformation; however, the correction must provide an equally compelling alternative explanation for the events.

How might the continued influence effect be mitigated in the initial example given above? Presumably, the initial misinformation suggesting food poisoning was continues to influence reasoning because of its causal coverage of the situation; that is, food poisoning explains why four related people of varying ages would all perish at the same time. In the absence of more information, it is difficult to generate viable alternatives that would also explain this event. However, another source did offer an alternative account with the correction:²

Monday: Reports family of four found dead in their home. The family had dined the night before at "China Garden."

Tuesday: Reports that food poisoning has been ruled out as family's cause of death. The deaths were blamed on a faulty furnace, which caused monoxide poisoning as the family slept.

Saturday: Surge in furnace service requests reported by local gas company.

The alternative (in italics) provides the same causal coverage of the family's deaths as the food poisoning explanation; namely, how four related individuals of differing ages would all succumb at the same time. The success of this correction is suggested by the increased requests for service as readers sought to check the safety of their own furnaces. When a compelling account can be offered as an alternative, people are less likely to fall back on the misinformation. Replacements are most successful when they account for the causal coverage of the events, regardless of their plausibility.

When no causal alternative is available, however, it appears to be very difficult to correct misinformation that plays a causal role in an account. When causal inferences are needed, people fall back on explanations they know are wrong in an attempt to "fill in" the causal gap. In a sense, the need for explanation may outweigh the known value of the information. A "good story" may not have to be true in order to be repeated and called upon when needed. Even when people are explicitly aware of the status of the misinformation, they may use it when necessary to explain events in the absence of an alternative account.

III. The Pragmatics of Correcting Misinformation

The evidence presented above suggests the continued influence effect is a major problem in successful correction. If a causal explanation is needed, attempts to avoid the use of misinformation depend on the presentation of an alternative cause. Ideally, corrections would not simply negate earlier information; instead, they would provide a further account that replaces the misinformation and provides an equally strong explanation for events. However, it is not always possible to provide an alternative account that is equally compelling: the true cause may be unknown. In such cases, people appear vulnerable to falling back on misinformation even though they know it has been discredited. In the absence of an alternative explanation, is there any way to mitigate the use of the discredited information? That is, can corrections be provided that will minimize use of the corrected information, even in the absence of a replacement?

The process of correction involves more than simply identifying previous information in memory and then negating it. Correction takes place in a context of communication that includes a network of conventions, such as offering only information believed to be true (maxim of quality), and relevant (maxim of relation) (Grice, 1989). Corrections may be particularly difficult to understand because they seem to violate these conversational conventions (Clark, Schreuder, & Buttrick, 1983). If the retraction is truthful, why would a well-intentioned source have presented the misinformation in the first place? And if the earlier information is believed to be accurate, how can the retraction be simultaneously accepted as accurate? Grice's conversational logic perspective suggests that corrections should pose problems in interpretation. Specifically, corrections are problematic when

the listener cannot believe that a coherent, “good faith” speaker would have uttered both the assertion and its correction. The listener must reconcile conflicting information in the correction; when unable to do so, they may fall back on the discredited information.

A. Conversational conventions

Grice's (1989) perspective asserts that the conventions of communication are important in interpreting the truth and relevance of statements. In particular, two maxims of communication are relevant to corrections.

1. **Relevance.** Listeners typically assume that conversational contributions will be relevant to the topic under discussion. In the case of corrections, the maxim of relation would imply that both the original assertion and its subsequent negation were relevant to the topic. But since the content of the two statements is directly contradictory, resolving the interpretation is problematic. A simple negation does not address the questions of the relevance of the initial misinformation. Because the maxim of relation implies that the information in the original assertion is indeed relevant to understanding the situation, listeners may continue to treat the information as relevant. This would cause the discredited information to continue to influence listeners' reasoning.

2. **Quality.** A parallel argument can be made for the impact of assumptions about informational quality upon understanding corrections. Listeners conventionally assume that speakers only provide information when they have good reason to believe it is true. In the case of a correction, the maxim of quality would imply that the speaker must have had good reason for believing that both the original assertion and its subsequent negation were true and accurate. A simple negation does not address the contradiction in the content of the two statements. Because the maxim of quality implies that the speaker had some good reason for believing the original assertion, listeners may continue to treat this information as accurate. Thus this information would continue to influence inferential understanding.

Several other studies have demonstrated the effects of conversational conventions on reasoning (Krosnick, Li, & Lehman, 1990; Schwarz, Strack, Hilton, & Naderer, 1991). This suggests that when corrections address only the literal content and not the conventions, questions may remain about how to interpret them.

B. Explanatory corrections.

To investigate the pragmatic implications of corrections, Julie Bush, Hollyn Johnson, and I compared several types of corrections intended to minimize the continued influence of misinformation. We again followed the news account paradigm in Wilkes and Leatherbarrow (1988) and Johnson and Seifert (1994). In one condition, the explain relevance correction, we explained why the original misinformation was no longer relevant to current conditions (“the closet had been empty because a trucker's strike had interfered with an anticipated delivery of paint and gas”). This correction explained why the information (that a delivery was expected) may have been initially presented. In an explain quality condition, we accounted for the contradiction by suggesting the

original information had, outside the speaker's awareness, been of poor informational quality (“the closet contained cans of coffee and soda rather than cans of paint and gas”). This correction explained how the presentation of misinformation may have occurred in error. We predicted that explanatory corrections would assist subjects in reconciling the pragmatic contradictions in the correction. If the explanations release subjects from assuming that the original assertion was accurate and relevant, they may show less influence from the misinformation than would subjects receiving a simple negation.

An alternative approach to improving the correction message may be to attempt to rule out alternative interpretations subjects may generate in their attempts to reconcile the correction. For example, subjects might resolve the contradiction between the misinformation and correction statements by inferring that volatile materials were present somewhere else on the premises, and so are still available play a role in the fire. This reasoning addresses the contradiction by combining the information and its negation in a novel inference not supported by the text. Enhancing the coverage of the correction statement may be helpful in ruling out alternative accounts not intended by the correction. An enhanced negation correction was tested to determine if ruling out this combination of messages helps subjects to adopt the correction as the valid conclusion. In the simple negation condition, the correction message told subjects that “the closet had actually been empty before the fire,” while the enhanced negation correction additionally asserted that “no paint or gas had ever been stored on the premises.” Elaborating on the literal content of a correction statement may make it easier to understand and avoid the “somewhere else” interpretation, but it may not help in resolving the contradiction. As a consequence, subjects receiving enhanced negation may be more uncertain about the validity of the correction, and show more influence from the discredited information. In particular, if the literal content of a statement is uninformative, its reiteration may raise doubt about the truth status of the statement (cf. Gruenfeld & Wyer, 1992; Wegner, Wenzlaff, Kerker, & Beattie, 1981; Wegner, Coulton, & Wenzlaff, 1985).

We also included two control conditions: a no-reference control, in which subjects never heard any suggestion that gas and paint had been involved, and an undisputed reference control, in which subjects heard without contradiction or correction that there had been paint and gas in the closet at the time of the fire. These conditions were intended to provide a baseline and a ceiling, respectively, for how often subjects report volatile materials as a factor in the fire. We predicted that subjects receiving either of the explanatory corrections (Explain Quality and Explain Relevance) would make fewer influenced inferences than those in the Simple Negation or Enhanced Negation conditions because the explanations would release these subjects from maintaining the original assertion. In addition, we expected that subjects receiving the enhanced negation correction would show a strong continued influence effect because it may backfire and raise doubt about the validity of the correction.

As in earlier studies, only a small number (8%) of subjects who saw a correction named the volatile materials as the “cause” of the fire, compared with 5% of the no-reference control subjects. There were also no memory differences among the groups for the reports. Figure 1 shows the mean

number of volatile materials inferences and direct references for subjects by condition. The undisputed reference condition resulted in significantly more volatile materials inferences than in the two negation conditions; the two negation conditions resulted in significantly more inferences than in the two explanatory correction conditions; and the two explanatory correction conditions resulted in significantly more inferences than in the no-reference condition. When the two explanation conditions are compared to the simple negation condition, the improvement in use of the misinformation is only marginally significant. This suggests the explanations are helpful in reducing use of misinformation, but not as effective as replacement with an alternative. As discussed above, when the negation is followed by a replacement cause, volatile materials inferences occur significantly less often compared to the negation condition.

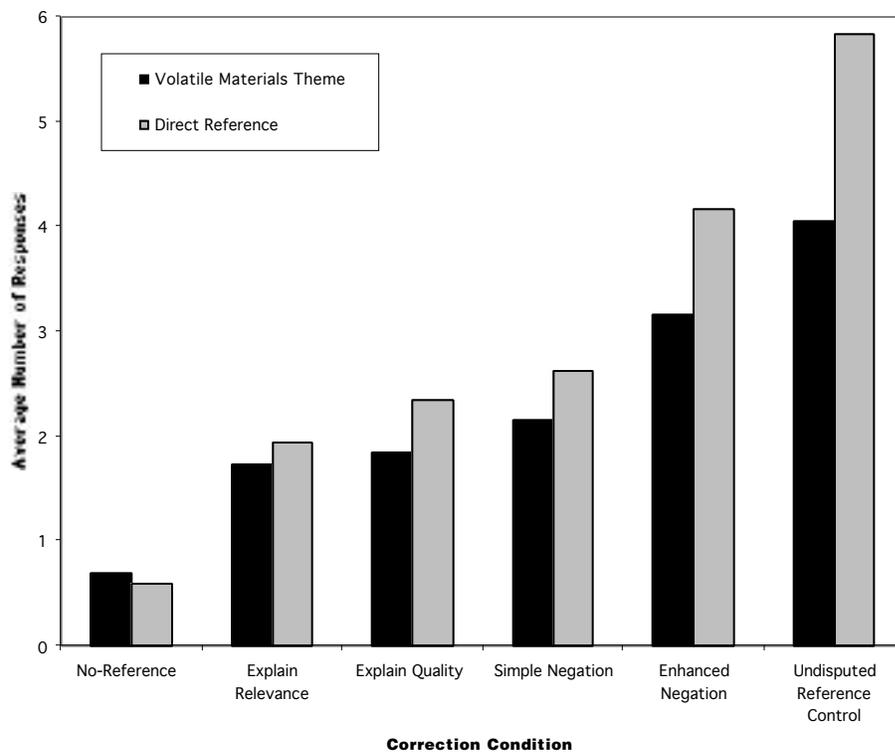


Figure 1. Mean Number of Volatile Materials Inferences and Direct References by Correction Condition.

The same pattern of results was obtained for the direct references measure: undisputed reference had significantly more volatile materials references than the two negation conditions, the two negation conditions had significantly more references than the two explanatory correction conditions, and the two explanatory correction conditions contained significantly more references than the no-reference condition. Once again, the continued influence effect was observed in most subjects: Only 23% failed to make at least one direct, uncontroverted reference to the volatile materials.

C. Qualifying inferences

These results provide support for the hypothesis that a correction is more effective if it contains an explanation addressing why the conflicting reports were given in the first place, and more generally, that conversational pragmatics are an important factor in successful corrections. The pragmatics of the correction statement may also affect whether and how subjects express uncertainty about the presence of the volatile materials. For example, a subject who said, "it seems that there was paint in the closet" would convey more uncertainty than one who said, "There was paint in the closet." To examine differences in how the volatile materials were referred to by subjects, we recoded the volatile materials measure as “unhedged” vs. “hedged” inferences. “Unhedged” excluded references to the volatile materials that were couched in terms of uncertainty (e.g., "if there was paint," "supposedly there were explosives in the closet"), or that suggested an integration of elements from both the original information and the correction (e.g., “gas was put back into the closet before the fire”).

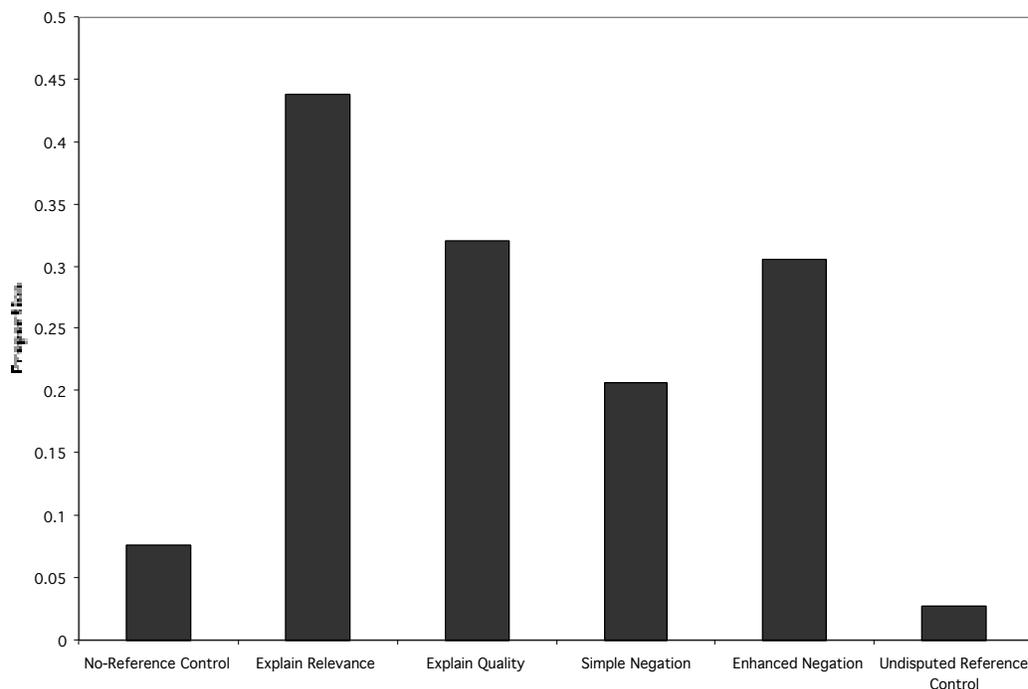


Figure 2. Mean Proportion of Hedged to Unhedged Volatile Materials Inferences by Correction Condition.

The proportion of “hedged” to “unhedged” statements for each condition is shown in Figure 2. All four correction groups used more hedges than the two control groups. The two explanation corrections resulted in a higher proportion of hedged responses than the simple negation group. Even the enhanced negation group had more hedges than the simple negation, perhaps reflecting a decreased belief in the correction. Compared to the three alternative corrections, the simple negation condition results in the most references to the volatile materials that are not qualified in some way.

This suggests simple negation may be less successful because it affords no suggestion about how to reconcile the contradiction. The two explanatory conditions can be considered successful in that they make no more unhedged inferences than the no-reference control group. That is, their use of the misinformation is rarely stated as if they believe it to be true (without qualification); instead, the inferences in the explanation conditions were qualified with some reference to uncertainty. However, in the simple and enhanced negation condition, unhedged inferences occurred significantly more often than in the no-reference control condition.

This analysis suggests that the pragmatics of a correction can affect how people think and talk about the discredited information. Differences in corrections are reflected in how subjects discuss the corrected information in their reasoning. For the explanatory corrections, subjects either qualified their inferences about the volatile materials (“If there were paint in the closet”), or rarely made use of the volatile materials (their “unhedged” inferences did not differ from the control group). For the negation conditions, only the simple negation showed a low rate of “hedged” inferences, but both simple and enhanced negation groups had significantly more unhedged inferences. By addressing the nature of the contraction between initial misinformation and correction, the Explain Relevance and Explain Quality corrections reduced the impact of the misinformation on later reasoning.

These results suggest that the continued influence of discredited information arises in part from the lack of explanation for the contradiction between the presentation of information and its later correction. People expect information presented to be relevant and informative; when correction occurs, these expectations are violated. When corrections can address these conversational expectations, they can mitigate one cause of the continued influence effect. While the explanatory correction groups showed less continued influence (than the negation conditions), they nevertheless made more inferences than did control subjects who never saw the misinformation. However, their inferences were also qualified with comments about the uncertain status of the misinformation. This is consistent with the idea that subjects faced with corrections are sensitive to explanations for why the conflict occurred. In the two explanatory conditions, subjects were offered a suggested explanation, and rarely directly attributed the cause to the volatile materials without including some qualification. These subjects continued to speak of the discredited information in ways that did not literally conflict with the discrediting statement.

An alternative approach to correction is to address any attempt to resolve the contradiction by positing the presence of the volatile materials someplace else on the premises. In the enhanced negation condition, the clarification added to rule out this possible interpretation backfired, and resulted in both more hedged and unhedged inferences compared to the simple negation group. If ambiguity in the correction alone were the cause of the continued influence of misinformation, the enhanced correction should have been more successful; however, adding information without resolving the contradiction did not reduce reliance on the misinformation.

The power of misinformation appears to arise from its legitimacy as initially presented. Direct negation of the misinformation, no matter how stated, does not address the initial value of the information as true and relevant, nor explain how the negation may have arisen. When explanations

are offered for the correction, subjects are less likely to use misinformation, and are careful to qualify their statements referring to it.

Further research is necessary to demonstrate the generality of these findings to other types of materials. The findings described here are based on variations of a single story account. It is important to establish that explanations of quality and relevance as pragmatic conventions have the same mitigating influence on other types of accounts. If so, corrections may be improved by address their inherent contradiction. Corrections are problematic because they require the integration of opposing information while maintaining belief in the truth and relevance of all of the information. Addressing these conversational pragmatics may be key in maintaining the coherence of an account, and the perceived value of the information presented.

IV. Effective Correction of Misinformation

Corrections are not always easy to understand, or to convey. When an assertion is made and then denied (e.g., "X is true ... actually, X is untrue"), a contradiction is generated. Consequently, corrections are likely to require additional processing time and effort to integrate into an account (Christianson, Hollingworth, Halliwell, & Feirarra, 2001). In addition, the memory representation of the account may be forever altered (Johnson & Seifert, 1998). Successful correction appears to require assisting the reader in resolving this contradiction. The findings discussed here suggest that the most powerful correction includes an alternative account to replace the misinformation (Johnson & Seifert, 1994). The alternative appears to be most effective when it can account for the causal features in the story that are left unexplained when omitting the misinformation (Johnson & Seifert, 1999). Another way to resolve contradiction is to address why the initial information may not be relevant now, or of good quality. Grice (1989) has argued that people expect speakers to offer only information that is truthful and conversationally relevant. Corrections that include explanations of why these conversational conventions were violated result in fewer references to the misinformation. Corrections that explain why misinformation was presented may release subjects from expecting the original assertion to be informative and relevant.

The empirical findings discussed here suggest that the continued influence of misinformation in memory is a powerful influence on later reasoning. Even when the correction is accurately reported, people still fall back on misinformation. In all of the studies described (Wilkes & Leatherbarrow, 1988; Johnson & Seifert, 1994; Wilkes & Reynolds, 1996; Johnson & Seifert, 1999; Bush, Johnson, & Seifert, 1994), the misinformation played a key role in the causal structure of the account. When asked questions, people appeared to fall back on misinformation in order to account for otherwise unexplained causal features. The negated information played a role in the causal account of events, and removing it left a "gap" that was unexplained. Such circumstances may be required for the continued influence effect to occur; if corrections involve material that is not causal in nature, or that is peripheral to story coherence, the continued influence of misinformation may not occur. Corrections that threaten information central to an account (van den Broek, 1990) and involved in

the causal structure of a story (Trabasso & van den Broek, 1985) are most likely to be affected by corrections as described here.

Similarly, when contradiction is expected, or when there are low expectations for information quality, misinformation may not influence later reasoning. The continued influence effect may depend upon the reader's assumptions about the validity of the information. Was the initial information or the correction untruthful? Was the information not relevant? How could both statements be true and relevant? With a simple negation of earlier information, the answer to these questions is not apparent. Readers are left to generate their own resolution of the conflict, which may not sufficiently address the status of the misinformation. The results described here suggest that modifying corrections to explicitly address the conventions of communication helps the reader to resolve the contradiction. However, if no contradiction is detected, or the reader perceives all of the information to be of low quality, misinformation may be successfully avoided in later reasoning.

These results suggest that misinformation can be propagated through later inferences despite even direct, immediate correction. The continued use of misinformation following correction may also provide evidence about the inference processes taking place during reading. Under a constructivist account (e.g., Trabasso & Suh, 1993), readers make global inferences wherever causal links are detected. Readers actively attempting to draw causal connections would notice the correction of potential causes, and attempt to reconcile the contradiction during reading. However, in our studies, subjects avoided addressing the implications of the correction until later questions forced the issue. This appears more consistent with a minimalist account (McKoon & Ratcliff, 1992), where readers make inferences needed to maintain local coherence, but avoid more global inferences while reading. In this account, the initial information could be tagged as "incorrect," but no further processing done to resolve the contradiction during reading. Then, when questioned, the reader is faced with the need to reconcile the accounts and determine what the correction means for the other information in the account. This postponement of the intensive reasoning required to resolve the contradiction may be functional in avoiding expensive processing that may not be required. Later, when questioned, more time and resources may be available to reason about the relationships between events and the role of the discredited information. However, as Wilkes and Reynolds (1999) note, the resolution of corrections will always be affected by processing capacity limitations.

These results seem to contradict those from traditional verbal-learning paradigms dealing with directed forgetting (e.g., Bjork & Woodward, 1973). Both paradigms address how people deal with irrelevant information presented in the context of relevant information. However, the studies here find continued influence from misinformation, in contrast to the robust forgetting effects found in much of the directed forgetting literature (e.g., Bjork & Woodward, 1973; MacLeod, 1975). The continued influence effect may depend upon having a coherent, causally related account where a single or minimal correction has a significant impact on the construal of meaning. In directed forgetting studies using sentences (Geiselman, 1974), as opposed to lists of unrelated words, subjects forget the forget-cued sentences less effectively when they are thematically related to remember-

cued sentences, or to each other. The interrelatedness of events in an account, and their causal relationships, may be key to the occurrence of the continued influence effect. Wilkes and Reynolds (1999) demonstrated that the number of causal relationships may not matter, but the existence of a causal relationship is sufficient to result in the continued influence of misinformation.

The current work also differs from that on post-event misinformation (e.g., Loftus, Miller, & Burns, 1978). In that paradigm, subjects witness a complete event, and then receive new, misleading information during testing. More critically, the post-event misinformation paradigm usually presents misleading information about a peripheral detail in the situation, with subjects left unaware that they have been misled; otherwise, subjects do detect and reject misinformation that blatantly contradicts their perceptions (Loftus, 1979). In contrast, the research described in this chapter shows continued influence from misinformation that is central to one's understanding of the event, and occurs even when subjects openly detect, accept, and accurately report the correction on memory measures. This paradigm suggests the introduction of contradictory information need not be subtle or tricky; in fact, direct negation can lead to misinformation effects. While a tremendous amount of research has addressed the construction of memory representations during reading (c.f. Gernsbacher, 1990), very little has addressed the reuse or return to those same representations over time. In addition to current use, representations persist in memory, and are recalled and reused in new circumstances. Presumably, they are also updated with new information as they are referred to again. Change in memory representations over time and use is a central feature of everyday reasoning, and our theories of cognition will need to address this challenging problem.

Of course, we do not radically fail to comprehend corrections every time we are exposed to them. In many instances, an explanation or an alternative account for a correction is readily apparent, commonplace, or familiar from experience. We are often able to incorporate information about the quality of the source (e.g., hearing it from a friend vs. watching a news report), changes in available knowledge (e.g., DNA-based exoneration of convicted criminals), and simple miscommunication (e.g., "I was referring to this case, not that one"). However, in settings where the quality of information is assumed to be high, more influence from corrections may occur. For example, in courtroom settings, the reasons for "striking" or disregarding information are often described in legal terms or left implicit. If the reason for the contradiction is not clear to jurors, they may struggle with continued influence from stricken information (Caretta & Moreland, 1983; Schul & Bernstein, 1985; Thompson, Fong, & Rosenhan, 1981). Similarly, in educational settings, information from instructors and texts are assumed to be of very high quality. When correction must occur, providing an explanation for its occurrence may reduce, but not eliminate, the continued influence effect. Intuitively, it seems helpful to present initial information accurately and avoid errors in instruction; however, little evidence exists about the impact of contentful corrections on learning (c.f. Seifert & Hutchins, 1992).

These findings are important because of the dynamic nature of information. Typically, reports about an event from many sources appear over long periods of time, and information continues to unfold about causes. Because any information available will be reported at the time, later reports

must frequently contradict initial information. For the recipient, corrections pose an inherent contradiction: How can one of these pieces of information be wrong? And if the misinformation is central to an account, how can it be omitted without replacing it? As a whole, these studies suggest the direct negation of information is not successful at “removing” misinformation from memory. Instead, people benefit from corrections that either replace the misinformation with a better alternative, or at the least, explain how the contradiction came to be. More generally, we might conclude that corrections are not solely informational, but are also interpreted within a pragmatic context of conventions about when and how information is communicated (Grice, 1989). As a consequence, successful correction of misinformation must address not only the informational content, but also the reader’s experience of contradictory information.

Footnotes

¹ Detroit’s WDIV TV News, November, 1991

² Ann Arbor News, November, 1991.

References

- Albrecht, J. E., & O'Brien, E. J. (1993). Updating a mental model: Maintaining both local and global coherence. Journal of Experimental Psychology: Learning, Memory, & Cognition, 19(5), 1061-1070.
- Anderson, C. A. (1982). Inoculation and counterexplanation: Debiasing techniques in the perseverance of social theories. Social Cognition, 1, 126-139.
- Anderson, C. A. (1983). Abstract and concrete data in the perseverance of social theories: When weak data lead to unshakeable beliefs. Journal of Experimental Social Psychology, 19, 93-108.
- Anderson, C. A., Lepper, M. R., & Ross, L. (1980). Perseverance of social theories: The role of explanation in the persistence of discredited information. Journal of Personality and Social Psychology, 39, 1037-1049.
- Anderson, C. A., New, B. L., & Speer, J. R. (1985). Argument availability as a mediator of social theory perseverance. Social Cognition, 3, 235-249.
- Baker, L., & Wagner, J. L. (1987). Evaluating information for truthfulness: The effects of logical subordination. Memory and Cognition, 15, 247-255.
- Belli, R. F. (1989). Influences of misleading postevent information: Misinformation interference and acceptance. Journal of Experimental Psychology: General, 118, 72-85.
- Bjork, R. A., & Woodward, A. E. (1973). Directed forgetting of individual words in free recall. Journal of Experimental Psychology, 99, 22-27.
- Bush, J. G., Johnson, H. M., & Seifert, C. M. (1994). The implications of corrections: Then why did you mention it? Proceedings of the Sixteenth Annual Cognitive Science Society Conference, Atlanta, Georgia, 112-117.
- Carretta, T. R., & Moreland, R. L. (1983). The direct and indirect effects of inadmissible evidence. Journal of Applied Social Psychology, 13, 291-309.
- Christianson K, Hollingworth, A., Halliwell, J.F., & Feirarra, F. (2001). Thematic roles assigned along the garden path linger. Cognitive Psychology, 42(4): 368-407.
- Clark, H. H., Schreuder, R., & Buttrick, S. (1983). Common ground and the understanding of demonstrative reference. Journal of Verbal Learning and Verbal Behavior, 22, 245-258.
- Conway, M., & Ross, M. (1984). Getting what you want by revising what you had. Journal of Personality and Social Psychology, 47, 738-748.
- Fleming, J., & Arrowood, A. J. (1979). Information processing and the perseverance of discredited self-perceptions. Personality & Social Psychology Bulletin, 5, 201-205.
- Geiselman, R. E. (1974). Positive forgetting of sentence material. Memory & Cognition, 2(4), 677-682.
- Gernsbacher, M. A. (1990). Language comprehension as structure building. Hillsdale, NJ: Lawrence Erlbaum Assoc.

Golding, J. M., Fowler, S. B., Long, D. L., & Latta, H. (1990). Instructions to disregard potentially useful information: The effects of pragmatics on evaluative judgments and recall. Journal of Memory and Language, *29*, 212-227.

Grice, H. P. (1989). Logic and conversation. In H. P. Grice (Ed.), Studies in the way of words (pp. 1-143). Cambridge, MA: Harvard University Press.

Gruenfeld, D. H., & Wyer, R. S. (1992). Semantics and pragmatics of social influence: How affirmations and denials affect beliefs in referent propositions. Journal of Personality and Social Psychology, *62*, 38-49.

Hertel, P. T. (1982). Remembering reactions and facts: The influence of subsequent information. Journal of Experimental Psychology: Learning, Memory, and Cognition, *8*, 513-529.

Jennings, D., Lepper, M. R., & Ross, L. (1981). Persistence of impressions of personal persuasiveness: Perseverance of erroneous self-assessments outside the debriefing paradigm. Personality & Social Psychology Bulletin, *7*, 257-263.

Johnson, H. M., & Seifert, C. M. (1994). Sources of the continued influence effect: When discredited information in memory affects later inferences. Journal of Experimental Psychology: Learning, Memory, and Cognition, *20*(6), 1420-1436.

Johnson, H. M., & Seifert, C. M. (1998). Updating accounts following a correction of misinformation. Journal of Experimental Psychology: Learning, Memory, & Cognition, *24*(6), 1483-1494.

Johnson, H. M., & Seifert, C. M. (1999). Modifying mental representations: Comprehending corrections. In S. R. Goldman & H. van Oostendorp (Eds.), The Construction of mental representations during reading (pp. 303-318). Mahwah, N.J.: Lawrence Erlbaum Asso.

Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. Psychological Review, *95*, 162-182.

Krosnick, J. A., Li, F., & Lehman, D. R. (1990). Conversational conventions, order of information acquisition, and the effect of base rates and individuating information on social judgment. Journal of Personality and Social Psychology, *59*, 1140-1152.

Loftus, E. F. (1979). Reactions to blatantly contradictory information. Memory and Cognition, *7*, 368-374.

Loftus, E. F., Miller, D. G., & Burns, H. J. (1978). Semantic integration of verbal information into a visual memory. Journal of Experimental Psychology: Human Learning and Memory, *4*, 19-31.

McCloskey, M., & Zaragoza, M. (1985). Misleading postevent information and memory for events: Arguments and evidence against memory impairment hypotheses. Journal of Experimental Psychology: General, *114*(1), 1-16.

MacLeod, C. M. (1975). Long-term recognition and recall following directed forgetting. Journal of Experimental Psychology: Human Learning & Memory, *104*(1), 271-279.

McKoon, G., & Ratcliff, R. (1992). Inference during reading. Psychological Review, *99*, 440-466.

Millis K K., & Erdman B. J. (1998). Comprehending news articles: Updating the news. Poetics, 25 (6): 343-361.

O'Brien, E.J., Rizzella, M.L., Albrecht, J.E., & Halleran, J.G. (1998). Updating a situation model: A memory-based text processing view. Journal of Experimental Psychology: Learning, Memory, and Cognition, 24(5),1200-1210.

Ross, L., Lepper, M. R., & Hubbard, M. (1975). Perseverance in self-perception and social perception: Biased attributional processes in the debriefing paradigm. Journal of Personality and Social Psychology, 32, 880-892.

Schul, Y., & Burnstein, E. (1985). When discounting fails: Conditions under which individuals use discredited information in making a judgment. Journal of Personality & Social Psychology, 49, 894-903.

Schwarz, N., Strack, F., Hilton, D. J., & Naderer, G. (1991). Judgmental biases and the logic of conversation: The contextual relevance of irrelevant information. Social Cognition, 9, 67-84.

Seifert, C. M., & Hutchins, E. L. (1992). Error as opportunity: Learning in a cooperative task. Human-Computer Interaction, 7(4), 409-435.

Thompson, W. C., Fong, G. T., & Rosenhan, D. L. (1981). Inadmissible evidence and juror verdicts. Journal of Personality and Social Psychology, 40, 453-463.

Trabasso, T., & Suh, S. (1993). Understanding text: Achieving explanatory coherence through on-line inferences and mental operations in working memory. Discourse Processes, 16, 3-34.

Trabasso, T., & Sperry, L. L. (1985). Causal relatedness and the importance of story events. Journal of Memory and Language, 24, 595-611.

Trabasso, T., & van den Broek, P. W. (1985). Causal thinking and the representation of narrative events. Journal of Memory and Language, 24, 612-630.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. Cognitive Psychology, 5, 207-232.

van den Broek, P. (1990). Causal inferences and the comprehension of narrative texts. In A. C. Graesser & G. Bower (eds.), The psychology of learning and motivation (Vol. 25): Inferences and text comprehension. San Diego: Academic Press, Inc.

Wegner, D. M., Coulton, G. F., & Wenzlaff, R. (1985). The transparency of denial: Briefing in the debriefing paradigm. Journal of Personality & Social Psychology, 49, 338-345.

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.

Wilkes, A. L. (1988). "Forget that. It was wrong:" Editing episodic memory to accommodate for error. Practical aspects of memory: Current research and issues, Vol. 1: Memory in everyday life (pp. 496-501). Oxford, England: John Wiley & Sons.

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

Wilkes, A. L., & Reynolds, D. J. (1999). On certain limitations accompanying readers' interpretations of corrections in episodic text. Quarterly Journal of Experimental Psychology: Human Experimental Psychology; 52A, 165-183.

Wyer, R. S., & Budesheim, T. L. (1987). Person memory and judgments: The impact of information that one is told to disregard. Journal of Personality and Social Psychology, 53, 14-29.

Wyer, R. S. & Srull, T. K. (1986). Human cognition in its social context. Psychological Review, 93, 322-359.

Wyer, R. S., & Unverzagt, W. H. (1985). Effects of instructions to disregard information on its subsequent recall and use in making judgments. Journal of Personality & Social Psychology, 48, 533-549.

Appendix A

Warehouse Fire Materials

Message 1: Jan. 25th 8:58 p.m. Alarm call received from premises of a wholesale stationery warehouse. Premises consist of offices, display room, and storage hall.

Message 2: A serious fire was reported in the storage hall, already out of control and requiring instant response. Fire engine dispatched at 9:00 p.m.

Message 3: The alarm was raised by the night security guard, who had smelled smoke and gone to investigate.

Message 4: Jan. 26th 4:00 a.m. Attending fire captain suggests that the fire was started by a short circuit in the wiring of a closet off the main storage hall. Police now investigating.

Message 5: 4:30 a.m. Message received from Police Investigator Lucas saying that cans of oil paint and pressurized gas cylinders had been present in the closet before the fire.

OR: (control): 4:30 a.m. Message received from Police Investigator Lucas saying that the closet was empty before the fire.

Message 6: Firefighters attending the scene report thick, oily smoke and sheets of flames hampering their efforts, and an intense heat that made the fire particularly difficult to bring under control.

Message 7: It has been learned that a number of explosions occurred during the blaze, which endangered firefighters in the vicinity. No fatalities were reported.

Message 8: Two firefighters are reported to have been taken to the hospital as a result of breathing toxic fumes that built up in the area in which they were working.

Message 9: A small fire had been discovered on the same premises, six months previously. It had been successfully tackled by the workers themselves.

Message 10: 10:00 a.m. The owner of the affected premises estimates that total damage will amount to hundreds of thousands of dollars, although the premises were insured.

Message 11: 10:40 a.m. A second message received from Police Investigator Lucas regarding the investigation into the fire. It stated that the closet reportedly containing cans of paint and gas cylinders had actually been empty before the fire.

OR (control): 10:40 a.m. A second message received from Police Investigator Lucas regarding the investigation into the fire. It stated that the two firefighters taken to the hospital had been released.

Message 12: The shipping supervisor has disclosed that the storage hall contained bales of paper; mailing and legal-size envelopes; scissors, pencils, and other school supplies; and a large number of photo-copying machines.

Message 13: 11:30 a.m. Attending fire captain reports that the fire is now out and that the storage hall has been completely gutted.

Appendix B Memory Questionnaire

Fact questions

1. What was the extent of the firm's premises?
2. Where did an attending firefighter think the fire started?
3. Where on the premises was the fire located?
4. What features of the fire were noted by the security guard?
5. What business was the firm in?
6. When was the fire engine dispatched?
7. What was in the storage hall?
8. What was the cost of the damage done?
9. How was it thought the fire started?
10. When was the fire eventually put out?

Inference questions

11. Why did the fire spread so quickly?
12. For what reason might an insurance claim be refused?
13. What was the possible cause of the toxic fumes?
14. What was the relevance of the closet?
15. What aspect of the fire might the police want to continue investigating?
16. Why do you think the fire was particularly intense?
17. What is the most likely cause of the fire that workers successfully put out earlier?
18. What could have caused the explosions?
19. Where was the probable location of the explosions?
20. Is there any evidence of careless management?

Manipulation check questions

21. What was the point of the second message from the police?
22. Were you aware of any corrections in the reports that you read?