Journal of Experimental Psychology:
Learning, Memory, and Cognition
1994, Vol. 20, No. 6, 1420-1436

0278-7393/94/$3.00

Sources of the Continued Influence Effect: When Misinformation in Memory Affects Later Inferences

Hollyn M. Johnson and Colleen M. Seifert

Several lines of research have found that information previously encoded into memory can influence inferences and judgments, even when more recent information discredits it. Previous theories have attributed this to difficulties in editing memory—failing to successfully trace out and alter inferences or explanations generated before a correction. However, in Experiments 1A and 1B, Ss who had received an immediate correction made as many inferences based on misinformation as Ss who had received the correction later in the account (and presumably had made more inferences requiring editing.) In a 2nd experiment, the availability (Tversky & Kahneman, 1973) of the misinformation within the comprehension context was tested. The results showed that Ss continued to make inferences involving discredited information when it afforded causal structure, but not when only incidentally mentioned or primed during an intervening task. Experiments 3A and 3B found that providing a plausible causal alternative, rather than simply negating misinformation, mitigated the effect. The findings suggest that misinformation can still influence inferences one generates after a correction has occurred; however, providing an alternative that replaces the causal structure it affords can reduce the effects of misinformation.

One may often learn "facts" about an event that later turn out to be false or unfounded. Ideally, one might want a correction to diminish or eliminate subsequent effects from such misinformation so that one understands and reasons about it like those not exposed to the misinformation. However, previous research (e.g., Carretta & Moreland, 1983; Ross, Lepper, & Hubbard, 1975; Wilkes & Leatherbarrow, 1988; Wyer & Budesheim, 1987) has found that discredited information can continue to influence reasoning and understanding, under some conditions, despite an instruction to disregard it. As a result of the prevalence of misinformation, it is important to investigate how one reasons about accounts containing corrections and what conditions could lead to a *continued influence effect.*

Some research has found influence effects, even when two simple explanations for them are ruled out: First, that subjects do not notice the discrediting, and second, that they fail to make the connection between the disregard instruction and the information it refers to. Much research (Carretta & Moreland, 1983; Ross et al., 1975; Wilkes & Leatherbarrow, 1988) has found that subjects do recall or acknowledge the disregard instruction when asked about it directly, suggesting that they have encoded it and can retrieve and potentially comply with it. Wilkes and Leatherbarrow also manipulated whether a correction directly stated what one was to disregard or required subjects to infer this, but they found no difference in the amount of influence from the misinformation. This suggests that influence can occur even when subjects have made the connection between the disregard instruction and the information it refers to.

Several theories have proposed that misinformation can influence judgments because one integrates it into a memory representation before a correction occurs. Hastie and Park (1986) argued that one makes an immediate personality assessment when one hears a trait or behavior attributed to a person, and others have proposed that one integrates new knowledge about a person into one's previous representation (Wyer & Srull, 1986). In this case, corrected information may continue to influence personality judgments because one may still retrieve and use assessments based on the initial information (Wyer & Budesheim, 1987; Wyer & Unverzagt, 1985). Similarly, belief perseverance research has proposed that the more prediscrediting elaboration or explanation one does the more influence the discredited information has (Anderson, Lepper, & Ross, 1980; Ross et al., 1975). These *prior processing* theories raise a couple of issues about what occurs when one must understand an account containing a correction.

One issue is whether misinformation can influence inferences made after a correction occurs, as well as those made before. When making on-line inferences (Graesser, 1981; Hastie & Park, 1986) while comprehending an account, one may involve misinformation in them before learning that the information is incorrect. Subjects may then fail to go back and successfully trace out and alter these inferences when the correction occurs. Thus, one might retrieve and report such unedited inferences from memory at a later assessment, rather than generating new ones. In this case, continued influence would reflect problems in editing what subjects already have in

Hollyn M. Johnson (now at Washington State University—Tricities) and Colleen M. Seifert, Department of Psychology, University of Michigan.

This work was supported by the Office of Naval Research under Contract N00014-91-J-1128 to the University of Michigan. We thank Jennifer Coe, Andrew Jeung, Srinika Narayan, and Christy Marshuetz for assistance in conducting the experiments, and extend special thanks to Julie Bush for help with materials. We also thank Hal Arkes, Brian Ross, A. L. Wilkes, Arthur Glenberg, and other, anonymous reviewers for comments on a draft of this article.

Correspondence concerning this article should be addressed to Colleen M. Seifert, Department of Psychology, University of Michigan, 330 Packard Road, Ann Arbor, Michigan 48104. Electronic mail may be sent to seifert@umich.edu.

memory. However, inferences showing influence from misinformation could also originate, in two different ways, *after* a correction occurs. First, one may still need to make on-line, backward inferences after the correction to link postcorrection information to what one read previously, thereby maintaining coherence in the account. Second, one may make postcorrection inferences in response to questions that arise after one has fully comprehended the account (on request). In both of these cases, continued influence could reflect subjects retrieving and using the misinformation itself to make new inferences. Previous work has not addressed whether the continued influence effect only stems from "unedited" precorrection inferences or may also occur through postcorrection inferencing.

A second issue is whether continued influence effects occur because misinformation is merely more available than other information, as a result of its having appeared within the understanding context. The mention of misinformation may activate key concepts in memory for subjects who hear it, whereas control subjects, who do not hear the target information, would not have these concepts activated. Encoding a correction later may not sufficiently reduce this activation, leaving these concepts more available than other alternatives. Tversky and Kahneman (1973) have proposed that information one can easily retrieve from memory, because it is more available, can bias frequency judgments. Similarly, misinformation may influence one's inferences about an account because the activated concepts seem more subjectively probable and lead subjects to produce plausible, easy inferences without evaluating them within the context of the account. In contrast, subjects who do not have these concepts activated would produce fewer inferences consistent with them.

On the other hand, continued influence from misinformation in an account may occur as a result of the causal role it plays, which provides a coherent structure to the event described. When making an inference, one may search one's representation of the account for an antecedent that satisfies causal constraints, such as operativity and sufficiency (Trabasso & van den Broek, 1985). One may retrieve a candidate antecedent that would meet these constraints, were it not discredited, and then may either fail to note the correction while making the inference or find the account hard to structure or extend without using the misinformation. Thus, one may retrieve this information and use it because it provides connections that structure the account, which might otherwise remain fragmented.

This causal explanation predicts a larger continued influence effect when a correction merely negates earlier information and does not provide any causal alternative. Previous work suggests that people rarely generate alternative interpretations spontaneously when a correction negates prior information (Anderson, 1982; Massad, Hubbard, & Newtson, 1979), which may leave only the misinformation to provide causal structure. Van den Broek (1990) has argued that accurate memory for text depends on constructing a connected causal chain. If a negation would seriously disrupt such a chain, then the need to maintain coherence may override the correction instruction. However, misinformation may have less influence when a correction presents an alternative that also provides a causal structure to the account. Such information could take over the misinformation's role in the account, so one could incorporate this alternative and still maintain a coherent understanding.

To investigate these hypotheses about how continued influence from misinformation arises, we used a fire investigation scenario from a text comprehension paradigm developed by Wilkes and Leatherbarrow (1988). In Experiments 1A and 1B, we examined whether misinformation influences inferences only as a result of problems in editing precorrection inferences stored in memory or whether postcorrection inferences can also show influence. These experiments manipulated whether the correction statement occurred several messages after the information it discredited (as in Wilkes and Leatherbarrow, 1988) or, instead, immediately following it, thus limiting the opportunities subjects had for making precorrection inferences based on the misinformation. In Experiment 2, we manipulated the role that target concepts played in an account to test whether mere mention and activation of concepts present in misinformation lead to influence effects. If these effects depend on the information's role in the account's structure, rather than its mere activation in memory, causally implicated concepts should promote more "influenced" inferences than concepts mentioned incidentally. Finally, in Experiments 3A and 3B, we examined whether subjects show less continued influence from misinformation when a correction presents a plausible causal alternative, rather than just negating the target information.

Experiments 1A and 1B

In Experiments 1A and 1B, we manipulated how soon a correction appeared after a piece of misinformation to test the "editing difficulties" explanation for the continued influence effect. Each group read a series of messages describing a warehouse fire, which were based on materials from Wilkes and Leatherbarrow (1988). In correction conditions, subjects read that a short circuit occurred near a closet reportedly containing volatile materials (cans of paint and pressurized gas cylinders), whereas control subjects read that the short circuit occurred near an empty closet. In a delayed correction condition, subjects read five additional messages before learning that the closet did not contain volatile materials. In an immediate correction group, the correction message directly followed the message about the volatile materials. For all conditions, the early messages in the series were written to limit opportunities for forward and backward inferences that would link information about the volatile materials to other contents in the account. Experiments 1A and 1B differed only in how the correction messages were worded.

If influence effects occur only when subjects make inferences based on the misinformation before the correction and then must edit them, one would expect more influence in delayed correction groups than in immediate correction groups. Subjects in the delayed conditions would have a window within which they could potentially make additional, strategic, causal and coherence-maintaining inferences before the correction. When the correction then occurs, they might not be able to track down and alter them successfully. Several theories have presumed that editing difficulties underlie continued influence effects (e.g., Ross et al., 1975; Wilkes & Leatherbarrow, 1988;

Wyer & Budesheim, 1987). In contrast, subjects in an immediate correction condition would have less opportunity to make precorrection inferences involving the volatile materials and less concrete information (e.g., features of the fire) to make volatile-materials inferences. This would limit precorrection inferences to those that subjects could randomly generate (Kintsch, 1988) on the basis of the volatile-materials message. In this case, one would not expect as much influence from misinformation when it is immediately corrected. Alternatively, if subjects use the misinformation to generate new, postcorrection inferences, one would expect influence even when subjects receive an immediate correction. In Experiments 1A and 1B, we tested the locus of the continued influence effect (whether subjects make new, postcorrection inferences based on misinformation), and do not address whether any of the inferences made arise as a result of on-line processes.

Method

Subjects. The subjects were University of Michigan undergraduates who received course credit in an introductory psychology class for participating. Subjects were run in groups of 8 to 10, in single sessions lasting approximately 50 min. We conducted the two experiments separately, with 32 subjects participating in Experiment 1A and 33 subjects in Experiment 1B.

Design. Experiments 1A and 1B each featured three groups: a delayed correction condition, in which five messages about a warehouse fire intervened between the misinformation and the correction; an immediate correction group, in which the correction message appeared immediately after the misinformation; and a no-reference control group, in which subjects did not receive the message that was later corrected in the other groups. The delayed correction and no-reference control groups replicated groups used in Wilkes and Leatherbarrow (1988).

Materials. The materials were modified versions of a series of reports used by Wilkes and Leatherbarrow (1988) that described the investigation of a warehouse fire. The series consisted of 13 individual messages, each 2 to 4 sentences long (see Appendix A). The messages were combined into a booklet, with 1 message per page. For the no-reference control groups, the 5th message in the series stated that a closet in the warehouse was empty, and later messages did not controvert this information. For the discrediting groups, the 5th message stated that the closet contained cans of oil paint and pressurized gas cylinders, and the correction appeared as a 2nd message from a police investigator stating that the previous message regarding the closet's contents was incorrect and that the closet was empty. For the immediate correction groups, the correction appeared as Message 6 (immediately after the information it discredited), whereas for the delayed correction groups, this statement appeared as Message 11 (allowing for 5 intervening messages). For the control group, Message 11 was a notice from the police investigator that several firefighters had been released from the hospital. For all three conditions, references to features of the fire were deleted from Messages 1 to 4 so they would not become involved in backward inferences once one had encountered Message 5. This would also limit subjects' ability to form a schematic or stereotypical view of the fire and thus could limit predictive, forward inferences.

The text of the correction messages from the two experiments is shown in Table 1. The correction message used in Experiment 1A essentially replicated that in Wilkes and Leatherbarrow (1988). However, previous research (Baker & Wagner, 1987) has found that subjects have more trouble detecting contradictions when presented in subordinate, rather than main, clauses. In our experiments, subjects could potentially misinterpret this phrasing as stating that the cans and cylinders are empty, but not the closet itself, which would contribute to the continued influence effects reported by Wilkes and Leatherbarrow. To avoid this problem, we used in Experiment 1B a correction that directly asserted that the closet contained no volatile materials, rather than embedding this information in a subordinate clause.

Table 1
Correction Messages Used in Experiments 1A and 1B

Message/experiment
Complex/Experiment 1a
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 oil paint and gas cylinders had actually been empty before the fire.
Direct/Experiment 1b
10:40 a.m. A second message received from Police Investigator Lucas regarding the investigation into the fire. It stated that there were no cans of oil paint or gas cylinders in the closet that had reportedly contained them; the closet had actually been empty before the fire.

Two memory tests were also prepared: first, a summary sheet, in which subjects wrote a free recall of the reports' contents and what they thought was responsible for the fire (what caused it to happen), and second, an open-ended questionnaire adapted from Wilkes and Leatherbarrow (1988). The latter included 10 questions on facts directly presented in the messages, 10 other questions requiring the subjects to make inferences about the event, and 2 final questions assessing whether subjects were aware of any correction or contradiction in the series (questions from the open-ended questionnaire are shown in Appendix B). All questions appeared in the same order for each subject, with all fact questions appearing before any inference questions to prevent the latter from introducing biases. The 2 contradiction questions always appeared after the inference questions. As an intervening task, subjects did a series of unrelated similarity ratings.

Procedure. Each subject received a booklet of reports and was instructed to read through it at his or her own pace, but they were not to go back and reread any of the messages. Subjects were also told that they would be asked to recall the information later. When each individual subject had finished reading, the experimenter collected his or her booklet of reports, and the subject received the summary sheet. When everyone had finished this test, all subjects worked on the intervening task for 10 min. After this time had elapsed, subjects received the open-ended questionnaire and were instructed to answer each question on the basis of their understanding of the reports. The experiment ended when all subjects had completed the questionnaire. We chose the instructions and order of tasks to replicate those of Wilkes and Leatherbarrow (1988).

Results

Scoring. A coder who was uninformed about the experimental conditions scored both the summary sheet and the open-ended questionnaire (including fact and inference questions). Three inference measures were computed: a *thematic-inference* measure, a *direct-reference* measure, and a *global-cause* measure. The thematic-inference measure consisted of negligence theme responses made on the 10 inference questions from the open-ended questionnaire. These were responses consistent

with believing that the warehouse had contained carelessly stored volatile materials that caused or contributed to the fire, as would be reasonable if the information about the volatile materials had not been discredited. This theme encompassed references using key words from the discredited message (e.g., *oil, paint, gas[es], cans,* or *cylinders*), mentions of the closet itself without indications that it was empty, and attributions of carelessness or negligence (see Table 2 for sample answers consistent with the negligence theme). One might expect subjects who received no information about the volatile materials to make some inferences consistent with a negligence theme, but most answers would likely reflect themes other than negligence, such as intentional setting of the fire, stored stationery at the warehouse, and the structure of the building.

The direct-reference measure of inferencing consisted of a count of all direct and uncontroverted references to the volatile materials themselves (paint cans or gas cylinders), including all measures: the free recall, in answer to fact questions, or in answer to the inference questions. A reference was considered controverted if the subject also stated that a later message had shown this to be incorrect or untrue, or the subject provided the gist of the correction message. Expressing uncertainty about the presence of the volatile materials alone was not considered a controverted reference. This direct-reference measure indicates the influence of discredited information occurring anywhere in subjects' protocols and provides an estimate that does not rely on "thematic" interpretations. Finally, a global-cause measure was determined by scoring whether answers to the fire-cause question from the summary sheet involved uncontroverted reference to the volatile materials.

Three recall measures were also computed. First, a *story recall* measure consisted of the number of noncorrection idea units that subjects reported in their free recall on the summary sheet. The series of story reports was broken down into component idea units by using an adaptation of procedures described in Kintsch (1974). A story unit was scored as recalled if the subject reproduced a recognizable portion of its content; otherwise the unit was scored as absent. Only messages common to all three conditions were scored; thus, the story recall measure does not include the message about the closet's contents (empty or storing volatile materials) or the message from the police investigator (correction or update on firefighters) because these did not appear in the story for all conditions. Second, a fact recall measure was computed by scoring accuracy on the 10 fact questions from the open-ended questionnaire. Finally, a correction recall measure was determined on the basis of whether subjects in the correction groups accurately referred to it in either the summary or the open-ended questionnaire.

Table 2
Sample Negligence Theme Answers to Inference Questions

Question	Answer
What was the possible cause of the toxic fumes?	Burning paint.
Why do you think the fire was particularly intense?	Oil fires are hard to put out.
For what reason might an insurance claim be refused?	Because flammable items were not kept in a safe place.

Inferences. The analyses reported below exclude idiosyncratic responses and items left blank, which made up 14.2% of the responses to the inference questions. Analyses of variance (ANOVAs) were done using the thematic-inference measure. For both experiments, Table 3 shows the mean number of negligence theme inferences per subject within each group. The results show a significant difference on this measure: $F(2, 31) = 13.30, p < .0001, MS_e = 41.01$; $F(2, 31) = 7.47, p < .002$, $MS_e = 14.96$, respectively. Planned comparisons revealed significant differences between the correction groups and the control group, with the control group making significantly fewer inferences consistent with the negligence theme: $t(29) = 5.00, p < .0001, SE = 1.31$; $t(29) = 3.84, p < .001, SE = 1.08$, respectively, but no significant difference between the two correction groups ($ts < 1.5$ in both experiments).

Table 3 shows the mean number of direct references to the stored volatile materials for each group, in both experiments. An ANOVA of this direct-reference measure of influenced inferences showed a significant difference in both experiments: $F(2, 31) = 11.71, p < .0001, MS_e = 35.63$; $F(2, 31) = 5.74, p < .008$, $MS_e = 12.98$, respectively, with the correction groups showing more influence than the no-reference control groups. In Experiment 1A, post hoc analyses using Tukey's honestly significant difference (HSD) procedure showed significant differences between the no-reference control group and each of the correction groups, but no differences between the two correction groups ($p < .007$). This pattern also occurred in Experiment 1B, with a significant difference ($p < .006$) between the immediate correction group and the no-reference control group, and a marginally significant difference ($p < .08$) between the control and the delayed correction group. Overall, 95% and 91% of subjects in the correction groups, in the two experiments respectively, made at least one direct and uncontroverted reference to the volatile materials. In the control groups for the two experiments, 27% and 30% of the subjects made reference to similar materials.

In contrast to the findings above, only 24% and 9.1% of subjects in correction groups, in the two experiments respectively, reported the volatile materials in response to the global-cause question from the summary sheet. The rest of the responses to this global-cause question were similar to those given by the no-reference control group, such as faulty wiring and arson, or were left blank.

Recall. There were no group differences on the story recall or the fact recall measures ($Fs < 1$ for each, in both experiments). The overall means for story recall were 14.3 and 12.2 units in Experiment 1A and 1B, respectively; subjects correctly recalled a mean of 9.2 and 8.3 facts, respectively. The correction recall measure showed high levels of recall in the correction groups in both experiments (100% of the delayed correction groups from both experiments and 90% and 91.7% of the immediate correction groups from the experiments, respectively). Analyses omitting subjects in the correction groups who did not report noticing the correction showed no difference from the results found when including all subjects.

Table 3
Number of References and Inferences Consistent With Misinformation, by Group, From Experiments 1A and 1B

Reference/inference	Experiment 1A			Experiment 1B		
	Delayed	Immediate	Control	Delayed	Immediate	Control
Direct references						
M	3.27	4.40	0.82	2.10	2.75	0.60
SD	1.56	2.01	1.66	1.20	1.96	1.87
Negligence inferences						
M	3.82	4.90	1.09	3.10	3.25	1.10
SD	1.78	1.60	1.87	1.37	1.60	1.20

Discussion

The findings that subjects show influence from misinformation, whether corrected early or late in the sequence, replicate and extend those of Wilkes and Leatherbarrow (1988) and are consistent with work on belief perseverance (Ross et al., 1975). The results of both Experiments 1A and 1B suggest that influence from misinformation can also occur as a result of postcorrection inferencing processes, rather than only when subjects may have generated the inferences before the correction. Subjects in the immediate correction group presumably made few precorrection inferences and thus had less need to edit their memories; this suggests that subjects may retrieve the misinformation itself to generate postcorrection inferences. This could occur either on-line, as one makes backward inferences that link postcorrection information to what one read previously or reconstructively, in answering questions that arise after one has fully comprehended the account. The latter view is consistent with previous research on reconstructive memory (Conway & Ross, 1984; Spiro, 1980), which had found that subjects, perhaps inappropriately, use more recent information to reconstruct earlier information and attitudes.

Memory differences cannot account for the continued influence effect found here: The groups did not differ in recall of the reports, and over 90% of those in the correction groups recalled the correction. Thus, these subjects had the raw materials available to avoid inferences influenced by the misinformation, yet they did not do so. Further, over 75% of the subjects in both experiments failed to report the volatile materials when directly asked what was responsible for the fire (the global-cause question). This suggests that results are not due to simple lack of belief in the correction. The results also show a continued influence effect whether the correction message was complexly worded, as in Experiment 1A, or directly worded, as in Experiment 1B. However, the correction groups in Experiment 1B made fewer negligence inferences ($M = 3.18$) than those that occurred in Experiment 1A ($M = 4.36$), suggesting that a more direct message had some mitigating effect.

One could argue that the editing hypothesis could still account for these results because the experiments do not provide direct evidence that subjects in the immediate correction condition did not make precorrection inferences. However, there are several reasons why this interpretation is not plausible. First, the early messages in the series did not contain references to features of the fire that the volatile materials could explain. Without these references, one would need to make few if any backward inferences on encountering the volatile-materials message. Thus, subjects in the immediate correction group would not be likely to have generated such precorrection inferences, which would require editing after the correction. Second, it is unlikely that subjects would make many forward inferences because the earlier messages just mentioned the existence of a fire. Van den Broek (1990) argued that readers are more likely to make forward inferences when necessary and sufficient causal conditions constrain those inferences; in this case, the mention of the fire gives few clues to its specific characteristics, and so one might not expect many predictions until more information came in. Third, the open-ended questionnaire asked about characteristics of the fire mentioned after Message 6 (postcorrection for the immediate correction group) and thus would be unlikely to tap and assess inferences based on precorrection information.

In interpreting these results, however, one cannot conclude that postcorrection inferencing accounts for all continued influence found in prior research. The editing hypothesis may still account for the continued influence that occurs in studies that actively promote inference or attribution making before a disregard instruction occurs (e.g., Anderson et al., 1980; Schul & Burnstein, 1985). Further, Experiments 1A and 1B did not assess whether the influenced inferences in the delayed correction condition reflect precorrection or postcorrection inferences; both could contribute to the influence observed. However, even if subjects make few precorrection inferences in this condition, as suggested by a minimal inferencing account (McKoon & Ratcliff, 1992), this would still support the main claim of these experiments: Postcorrection inferencing involving misinformation can account for some of the continued influence effect observed.

Thus, these findings address the locus of the continued influence effect and suggest that it can occur as a result of processes operating after the correction, as well as those occurring before. One can also ask whether misinformation must originally play a causal role in the account to influence later reasoning. Previous research on comprehension (Graesser & Clark, 1985; Haviland & Clark, 1974) has suggested that causal relations give an account structure and provide a basis for further inferences. Subjects may have trouble finding an alternative way to structure an account coherently when causally central information is discredited, so they may continue to rely on it and use it in later inferencing. Some research

in belief perseverance (Anderson et al., 1980; Fleming & Arrowood, 1979; Ross et al., 1975) has proposed that misinformation has more influence when subjects try to explain reasons for their performance, which would presumably involve the misinformation in a causal structure; however, other work has not reported this explanation effect (Anderson, 1982, 1983; Jennings, Lepper, & Ross, 1981).

Alternatively, perhaps the mere mention of misinformation activates concepts that suggest plausible inferences. Prior work in decision making (Tversky & Kahneman, 1973) has proposed that more readily available information in memory can bias judgments. If a similar activation hypothesis can account for the continued influence of misinformation in accounts of events, this would suggest that any mention of the information, regardless of contextual support, could show similar influence. It would also suggest that the continued influence effect may be an artifact of the assessment, in some cases. If activation is solely responsible, one would expect influence effects to diminish over time, as activation lessens for key concepts introduced in the misinformation. In Experiment 2, we tested whether information merely mentioned in an account influences one's inferences or whether the information must also afford causal structure for it.

Experiment 2

In Experiment 2, we tested whether the continued influence effect depends on the sheer availability of the misinformation in memory, regardless of its role in the account. First, in a corrected causal reference group, subjects read that a short circuit occurred near a closet reportedly containing volatile materials. Here causal inferences are likely because conditions of operativity and sufficiency (Trabasso & van den Broek, 1985) are met. Subjects later learned that the closet did not contain volatile materials. In a second causal reference group, subjects read the same volatile-materials information as in the corrected causal reference group, but did not receive a correction later. The third condition was a noncausal reference group, which presented the volatile materials within the context of the account but as products in a store across the street from the fire location. In this condition, the information does not meet conditions conducive to involvement in causal inferences, such as operativity and sufficiency (Trabasso & van den Broek, 1985). The presence of volatile materials in a store across the street is not sufficient, in the context of the reports, to account for the warehouse fire; one would need to make many inferences beyond the story context (e.g., that someone moved the materials into the warehouse, that someone intended this, and so forth) to establish these materials in a causal role. For this group also, the series of reports did not include a correction. Finally, in a fourth, external reference condition, subjects read a series of reports that did not present information about the volatile materials or a correction. However, in an intervening task, they generated associates to key concepts introduced as misinformation in the corrected causal reference group. Previous research on implicit memory (Schacter, 1987) has found that processing a word list can increase the probability that subjects will produce those words later on in an ostensibly unrelated task (here, the postcorrection assessment).

If the causal structure of information affects its influence on inferences, one would expect little influence in the external and noncausal reference groups, in which the volatile materials information contributes little to one's causal understanding of the account. One would expect greater influence in the causal and corrected causal reference groups. At the other extreme, if the continued influence effect results simply from activation of plausibly related concepts in memory, one would expect the external reference group and the corrected causal reference group to show similar influence on inferences reported in the postcorrection assessment, even though the activation of the volatile material key words originates outside the story context for the former. Finally, if one considers the postcorrection assessment as an explicit-memory test (Schacter, 1987), subjects may only rely on activation of concepts that occur within the episodic representation of the report series. Thus, if the influence effect is due to activation of plausibly related concepts within a particular episodic context, one would expect few influenced inferences in the external reference group, but a higher number of influenced inferences in both the noncausal and the corrected causal reference groups.

Method

Subjects. Sixty-one University of Michigan undergraduates participated as partial fulfillment of an introductory psychology requirement. Subjects were run in single, 45-min sessions, in groups of 8 to 10.

Materials. The materials were modified versions of the report series used in the previous experiments. For the causal reference and the corrected causal reference groups, the fifth message stated that the closet contained volatile materials. The correction appeared as Message 11 for the corrected causal reference group; in the causal reference group, this message stated that several firefighters had been released from the hospital. In the noncausal reference group, Message 3 stated that the fire department was called by the manager of a store located across from the warehouse, which sold oil paint and pressurized gas cylinders. In this group, Message 5 stated that the closet in the warehouse was empty, and Message 11 was the same as in the causal reference group. Thus, the noncausal reference condition mentioned the volatile materials within the story context, but as a peripheral fact. Finally, the external reference group used the same series of reports as in the causal reference group, except that Message 5 stated that the warehouse closet was empty rather than containing volatile materials.

An intervening task was also constructed, made up of 102 pairings of category names, such as *city,* with initial letters, such as *F.* For each pairing, subjects were to think of a category exemplar that begins with that letter. Two versions of this task were used. Subjects in the external reference group saw the priming version, which presented the words *paint, gas,* and *container* as categories and also presented categories for which *paint, gasoline,* and *oil* are typical examples. Thus, by reading and generating these items, subjects in the external reference condition would activate key concepts introduced in the misinformation that the corrected causal reference group received. These items appeared among the 20 items presented on the first page of the task. Subjects in the other three groups received a second control version of the task, which was the same except that neutral items replaced those involving the key concepts from the misinformation message.

The two memory tests were the same as in the previous experiments. For the noncausal reference group, the open-ended questionnaire contained an additional question about the contents of the store across the street, appearing after the two contradiction questions.

Procedure. We used the same procedure as in the previous experiments.

Results

All dependent variables were coded as in the previous experiments.

Inferences. The following analyses exclude idiosyncratic responses and inference questions left blank (18%).

Table 4 shows the mean number of negligence thematic inferences and of direct references to the volatile materials. An ANOVA of the number of negligence thematic inferences showed a significant difference: $F(3, 57) = 9.16, p < .0001$, $MS_e = 26.34$. A post hoc analysis using Tukey's HSD procedure found no significant difference between the two causal reference groups or between the external and noncausal reference groups. However, both of the causal reference groups significantly differed from the other two groups ($p < .05$).

An ANOVA of the direct-reference measure, with group as a variable, also showed a significant difference, $F(3, 57) = 30.34, p < .0001, MS_e = 62.77$. A post hoc comparison using Tukey's HSD procedure again showed no significant difference between the two causal reference groups, which heard about the volatile materials, or between the two other groups that did not. However, both of the causal reference groups again significantly differed from the noncausal and the external reference groups ($p < .001$). Overall, 97% of the subjects in the two causal reference groups made at least one direct and uncontroverted reference to the volatile materials, versus 30% of the subjects in the noncausal and external reference groups.

Analysis of the global-cause question from the summary sheet showed that 22.6% of the subjects in the two causal reference groups attributed the fire to the volatile materials when directly asked what was responsible for it. In comparison, only 3% of subjects from the external–noncausal reference groups made such attributions. The rest of the responses attributed the fire to such causes as faulty wiring or arson or were left blank.

Recall. As in Experiments 1A and 1B, the groups did not significantly differ on the fact recall measure ($F < 1$), with an overall mean of 8.4 facts recalled correctly, or on the story recall measure ($F < 1.2$), with an overall mean of 13.02 units recalled. In the corrected causal reference group, 80% of the subjects recalled the correction. Analyses omitting subjects in this group who did not report noticing the correction showed no difference from the results found when including all subjects.

Other manipulation measures. In the noncausal reference condition, 80% of the subjects correctly recalled the incidental information about the store's contents. Analyses omitting subjects who did not recall this showed no difference from results found when including all subjects.

A coder who was uninformed about the experimental conditions scored the intervening task. Subjects received a point for each associate written in response to one of the three volatile-materials words presented (*oil, paint,* and *gas*) and for each volatile-materials key concept generated as an associate to a presented item. The results show that the external reference group made an average of 2.5 such responses. Among the other groups (who were not presented with any of these target items), only 2 subjects spontaneously generated such a response.

Table 4
Number of Inferences and References Consistent With Misinformation, by Group, From Experiment 2

Inference/ reference	Group			
	Causal reference	Corrected causal	Noncausal reference	External reference
Negligence inferences				
M	3.88	3.40	1.13	1.73
SD	1.96	1.88	1.19	1.62
References to volatile materials				
M	4.44	3.47	0.73	0.33
SD	1.71	1.81	1.22	0.72

Discussion

The results of Experiment 2 show that misinformation is more likely to influence inferences when it affords causal structure for the account and not simply when its key concepts remain activated and are therefore easily retrievable. Subjects in the external and noncausal reference groups made few inferences consistent with or mentioning the volatile materials. However, subjects in the external reference group processed and generated the key concepts in the intervening task, and 80% of the subjects in the noncausal reference group recalled the incidental information about the store's contents when directly questioned about it. These results suggest that both groups had the key concepts available, but that they failed to use them in making inferences, as a result of the lack of causal relevance. In contrast, subjects showed a substantial number of inferences involving the volatile materials in both the groups in which this information was mentioned within a causal context. This suggests that the causal structure that the information affords, and not its mere availability, can lead to influence on inferences about an account, even when a correction appears later.

There are a couple of potential explanations for why the present results do not replicate those of previous research (Golding, Fowler, Long, & Latta, 1990; Wyer & Budesheim, 1987; Wyer & Unverzagt, 1985), which has found that a disregard instruction at least partially mitigates effects of mistaken information. First, a majority of those studies have used an impression formation task, rather than a text comprehension task, and this may involve a different discounting procedure. Wyer and Budesheim (1987) proposed that people may have theories on how misinformation could influence their personality judgments and may use them to adjust their ratings. Subjects in our study may not have a comparable theory on how misinformation about the event would become incorporated into their inferences, and so they may not be able to adjust their responses accordingly. Second, subjects in the Wyer & Budesheim study knew their goal was to form an

impression, which they would base judgments on later. In contrast, in the present experiment, subjects were given recall instructions before the task, but no explicit instructions about comprehension. Thus, subjects in the person impression task may process the correction differently because they can monitor their discounting efforts in terms of that task. Without the ability to monitor their discounting efforts, our subjects' memories may contain blatantly contradictory beliefs about the event, which they might not realize until something draws their attention to it.

Alternatively, previous research (Golding et al., 1990; Wyer & Budesheim, 1987; Wyer & Unverzagt, 1985) also differs from what is reported in this article in how they motivated the disregard instruction. In previous work, subjects received a disregard instruction because the experimenter mistakenly read information about a completely different entity (a different person or work contract). Subjects showed more mitigation of the misinformation's influence in mistake conditions than when simply told to disregard information because it was confidential (Golding et al., 1990; Wyer & Unverzagt, 1985). In the present research, the correction negated information appearing earlier but did not address why the mistake occurred. Thus, subjects do not know why someone volunteered the misinformation in the first place, when providing it would violate conversational norms (Grice, 1975) of quality and relevance. They may accept the correction, in a limited or literal sense, but still believe that there is something to the original information, or else the reports would not have mentioned it. For example, subjects might accept the correction's literal content and believe that the closet was empty before the fire, but they might still believe that the volatile materials existed somewhere else or at some other time. In this case, the continued influence effects may reflect subjects' reasonable attempts to make sense of why the report originally mentioned the misinformation and to address the error's implications. Such a process would still result in comprehension differences between subjects who encounter misinformation and those who do not.

Overall, the present results provide stronger evidence that subjects come to understand the same event differently on the basis of whether an account contains a correction. Such findings would have implications for one's understanding of events in the world. Often, as one obtains more information about events over time, many initial speculations may turn out to be unfounded. However, even explicitly stated corrections often merely indicate that certain previous information is incorrect, and require the listener to reason out the implications of the change, or why someone provided what turned out to be misinformation. These experiments demonstrate that one cannot achieve complete discrediting, if that is one's intent, by merely negating the literal content of information. However, there may be other ways to provide corrections that would lead subjects to accept the initial report as a meaningless error (that the materials were not in the closet, or anywhere else in the warehouse, at any time).

Our findings are also consistent with other research that has examined similar effects, despite differing definitions of availability and activation. In research on belief persistence, Anderson, New, and Speer (1985) proposed that availability contributes to influence from discredited information. However, that work manipulated whether one had alternatives to the discredited information available, but it did not make claims about, or directly manipulate, the relative activation level of the discredited information itself. Subjects were also likely to have made connections linking the information to other world knowledge before it was discredited; this would more closely parallel the causal reference groups in this experiment than a mere-mention condition. Thus, our results are consistent with this work. Other previous research, on co-reference (O'Brien & Albrecht, 1991; O'Brien & Myers, 1987), has found that it is easier to reinstantiate a concept that has causal connections to other text material than one that is less causally implicated. Thus, causal connections may lead to more activation than a mere mention would, and therefore result in more influence. However, this is not inconsistent with the findings here, as it still supports the idea that both the causal and the corrected causal reference groups show influence due to the causal structure that the volatile-materials information affords.

On the basis of the causal hypothesis, one might also expect subjects to show less influence from misinformation when a correction provides a causal alternative, rather than simply negating earlier information. When a correction merely negates information, many aspects of the event are left unaccounted for. One may make inferences, and therefore responses, that maintain a coherent story versus having an account that contains many loose ends. However, often a correction provides an alternate fact, for example, "No, it wasn't a blue car, it was a red one." In this case, one has a concretely stated alternative that one can believe and understand in the same causal role that the misinformation held. Thus, one would not need to maintain the misinformation to have a coherent account. The alternative could allow one to, in essence, "drop" the misinformation from one's representation of the event or could compete and win out as a causal antecedent for use in inferences. In Experiments 3A and 3B we examined inferences made when a correction provides a new causal alternative along with a negation by using different story scenarios.

Experiment 3A

To determine whether misinformation has less influence when a correction provides a causal alternative, we used in Experiment 3A modified versions of the materials from the earlier experiments. For two groups, the series of reports implied an accident theme, which encompassed both the negligence theme from the earlier experiments and a theme involving the stationery supplies stored in the building. For two other groups, the reports implied an arson theme. The variable of implied theme was crossed with whether a correction occurred or not. Thus, the four groups were: (a) a negation group, as in previous experiments, in which a correction stated that the earlier information about volatile materials was not true, leaving one with a stationery supplies subtheme if one took the negation into account; (b) a mention control group, in which the reports presented information about the volatile materials and did not correct this later, leaving one with the negligence subtheme; (c) an arson control group, in which the

reports made no mention of the volatile materials (paint and gas) but instead mentioned that suspicious rags and steel drums (arson materials) were found in another location; and (d) an alternative group, in which subjects initially read about the volatile materials stored in the closet and later learned that this was not true, but arson materials had been found elsewhere. Throughout the rest of the article, *volatile materials* refers exclusively to the oil paint and gas cylinders reportedly stored in the closet, and *arson materials* refers to the gasoline-soaked rags and emptied steel drums used to intentionally set the fire.

If subjects can discredit information when the correction presents an alternative instead of just a negation, one would expect the alternative group to show fewer negligence inferences, and fewer references to the volatile materials, than the negation group would; at best, it would be no different from the arson control group, which never heard mention of the volatile materials. Subjects could create and report a coherent account of the event by relying on the information about the arson materials and thus would not use the information about the volatile materials in their inferences. However, if having a causal alternative does not limit the influence from misinformation, in the worst case, one would expect the alternative group and the mention control group to show a similar number of negligence inferences and references to volatile materials. To validate the arson manipulation, we performed comparisons of the arson control and mention control groups, which will show whether the implied themes are distinct enough to warrant interpretation of the other conditions or whether mention of these arson materials leads to responses involving the volatile materials by association, even though the reports do not directly mention them.

Method

Subjects. Eighty-one University of Michigan undergraduates participated as partial fulfillment of an introductory psychology requirement. Subjects were run in single, 45-min sessions, in groups of 10 to 12.

Design. The design was a 2 × 2 between-subjects factorial, with correction (yes vs. no) and implied theme (accident vs. arson) as variables. The mention control and negation conditions systematically replicated the causal reference group of Experiment 2 and the negation groups in the previous experiments, respectively. Two other groups read reports that suggested arson as the fire's cause. The arson control group received only information suggesting arson, and no contradictory information, whereas the alternative group first read about the volatile materials carelessly stored in a closet, learned this was an error, and then received information suggesting arson.

Materials. The materials were essentially the same as those used in the previous experiments, except that the series of reports contained 15 messages, with 2 new messages added to allow for further inferences that could discriminate between the negligence and arson themes. Appendix C shows the text of the new messages; Table 5 indicates how messages were combined to manipulate theme and correction in the different conditions.

Message 6 either read that a storage closet near where the fire started contained cans of paint and gas cylinders or read that it was empty. Message 13 either stated that the previous message regarding the closet's contents was incorrect and the closet was empty or stated that two firefighters had been released from the hospital. Message 14 either gave a list of supplies normally stocked by the warehouse or gave indications that gasoline-soaked rags had been found in suspicious circumstances. The summary sheet and open-ended questionnaire were the same as in the previous experiments, except that the questionnaire contained several additional questions, on topics like motives of the owner and the cause of a previous fire on the premises, to assess inferences consistent with the different themes used here.

Table 5
Content of Critical Messages by Group From Experiment 3A

Group	Message 6	Message 13	Message 14
Mention control	Volatile materials	Firefighters	Warehouse supplies
Negation	Volatile materials	Empty closet	Warehouse supplies
Alternative	Volatile materials	Empty closet	Arson materials
Arson control	Empty closet	Firefighters	Arson materials

Procedure. We used the same procedure as in the previous experiments.

Results

Scoring. The dependent variables used in the previous experiments were coded as described in Experiment 1. As before, the negligence thematic-inference measure included responses that explicitly used key words from the message about volatile materials in the closet (e.g., oil, paint, gas[es], cans, or cylinders). In addition, a coder who was uninformed about the experimental conditions scored the inference questions for responses consistent with an arson theme. This arson theme encompassed responses that explicitly used key words from the message presenting the arson materials (e.g., gasoline, rags, drums, or barrels) or expresses arson intent or motives for arson. The strict coding of key words from the original "volatile-materials" message versus the later "arson materials" message allowed the coding scheme to distinguish between the volatile and arson materials and themes. Answers containing a feature from both the arson and the negligence theme were coded as a "mixed reference." The following analyses do not include idiosyncratic responses or inference questions left blank (20.4%).

Inferences. In 2 × 2 ANOVAs, we tested whether the groups differed in number of thematic inferences consistent with either negligence and arson (see Table 6). There was a significant main effect of implied theme for both the negligence and the arson themes: for negligence, $F(1, 76) = 15.02$, $p < .0001$, $MS_e = 48.81$; for arson, $F(1, 76) = 27.98$, $p < .0001$, $MS_e = 95.90$. For both of these analyses, the main effect of correction and the interaction of correction and implied theme were not significant ($F < 1$). The two control conditions differed significantly on these two measures in planned comparisons: for negligence, $t(76) = 3.29$, $p < .002$, $SE = .57$; for arson, $t(76) = 3.54$, $p < .001$, $SE = .59$. Overall, the mention control group made few arson inferences and the arson control group made few negligence inferences.

To assess specific influences of the discredited information, we compared the direct-reference measure in a 2 × 2 ANOVA, with implied theme (accident vs. arson) and correction (yes vs.

Table 6
Number of Thematic Inferences and References Consistent With Misinformation, by Group, From Experiment 3A

Reference/ inference	Mention		Arson	
	Negation	Control	Alternative	Control
Negligence inference				
M	4.14	4.32	2.60	2.10
SD	2.17	2.40	1.50	1.41
Arson inference				
M	1.33	1.47	4.00	4.33
SD	1.32	1.84	1.69	2.54
Direct references				
M	3.71	3.47	3.25	0.29
SD	2.43	2.01	1.89	0.90

no) as the variables. This measure showed a significant interaction between groups, $F(1, 76) = 9.51, p < .003, MS_e = 33.45$ (see Table 6). A post hoc comparison using Tukey's HSD procedure showed a significant difference between the arson control group, which never heard about those materials, and each of the three groups that did ($p < .001$). In the three groups that heard mention of the volatile materials, 91.5% of the subjects made one or more direct and uncontroverted references to them, whereas in the arson control group, only 14.3% made spontaneous reference to similar materials.

Responses to the fire-cause question from the summary sheet were also scored as before in a global-cause measure. In the mention control group, in which information about volatile materials was not controverted later, 78.9% of the subjects attributed the fire to those materials. In the negation and alternative groups, in which information about the volatile materials was discredited, 28.8% and 22.2% of the subjects attributed the fire to those materials, respectively. On the other hand, subjects in the alternative and arson control groups attributed the fire to arson (77.8% and 81.0%, respectively).

Recall. As in the prior experiments, the four groups did not differ on story recall or in fact recall measures ($F < 1.7$ and $F < 2.61$, respectively). The overall mean for story recall was 11.19 units, and subjects recalled a mean of 8.2 facts. On the correction recall measure, both of the correction groups showed high levels of recall of the correction, with 95.2% of the negation group and 75.0% of the alternative group recalling it. Analyses omitting subjects in the correction groups who did not report noticing the change showed no difference from the results found with all subjects.

Discussion

The thematic inference results show that misinformation has less influence when subjects receive a correction that provides an alternate causal explanation, rather than one that only negates the information. Subjects in the alternative group made as many arson inferences as a control group for whom arson was the only theme given, and the two groups did not differ in negligence inferences, despite the alternative group having heard mention of the volatile materials. However, subjects in the negation group made as many negligence inferences as a control group for whom the information on volatile materials was valid. These findings also support the claim that the continued influence effect depends on the causal structure that the misinformation affords an account and are consistent with proposals that one may reason by using discredited information when other alternatives do not come to mind (Anderson et al., 1985; Wegner, Coulton, & Wenzlaff, 1985).

The results also suggest that a correction accompanied by a causal alternative works better than one entailing only negation, at least when the alternative adequately fills the causal role the misinformation would hold. This could occur for several reasons. First, misinformation that is only negated may interfere with subjects' ability to generate alternatives, in effect fixing them in one interpretation of the event. Providing an alternative may break this set and allow subjects to consider other interpretations. Second, the alternative and the misinformation in memory may compete when one must make an inference. If both pieces of information come to mind, one might have to deliberate between the two, which could make one more likely to heed the correction. Finally, the alternative may help one explain why the misinformation was mentioned in the first place. In this case, subjects may reason that someone volunteered the original information about volatile materials because they had mistaken them for the arson materials, the presence of which was confirmed later. With a simple negation, however, subjects do not know how the misinformation originated and so may infer that the materials must exist somewhere to explain their mention.

These results also replicate those in the previous experiments and provide more evidence that continued influence effects can occur despite subjects noticing and accepting the correction in some sense. First, over 90% of subjects in the negation group recalled the correction, suggesting that the effect is not just due to subjects failing to detect a discrepancy in the account. Second, only 29% of the subjects in the negation group mentioned the volatile materials in response to the fire-cause question, as compared with 79% in the mention control group, in which the information was valid. One would not expect such a large difference if subjects simply did not believe the correction. However, over 90% of the subjects in both of these groups made at least one uncontroverted reference to the volatile materials, which provides evidence that the effect is not merely due to responses from subjects who did not accept the correction.

However, a causal alternative does not necessarily prevent all influence from misinformation. The results also show that subjects made more direct references to the volatile materials if they were mentioned at all, regardless of the type, or existence, of a correction. This may occur because a correction does not expunge the discredited information from memory, so one still has some chance of retrieving it. It may then beat out the alternative, or provide a better or more plausible explanation, so one may involve it, rather than information about arson materials, in one's inferences. Or, subjects in the alternative condition may be more likely, in this case, to link the volatile materials and the arson materials and see them as being the same thing, thus interpreting the earlier message as correcting the location of the volatile materials and not

whether those materials existed at all. This might be especially likely because the two alternatives share a causal role and many of the properties that allow them to fill that role. Some support for this option comes from the finding that subjects in the alternative group showed lower recall of the correction (75%) than did the negation group (over 90%).

Experiment 3B

One could argue that the continued influence effects found in the previous experiments are an artifact of the particular warehouse fire scenario used, and thus that findings on how to mitigate the effect might not generalize to other stories. A number of studies have found similar continued influence effects by using different story materials in comprehension tasks (Hertel, 1982; Spiro, 1980; Wilkes & Leatherbarrow, 1988) and in a variety of tasks other than comprehension: Subjects have been found to involve discredited information in making judgments about task ability in others and oneself (Ross et al., 1975), social theories (Anderson et al., 1980), person impressions (Schul & Burnstein, 1985; Wyer & Budesheim, 1987; Wyer & Unverzagt, 1985), decisions about likely activities (Schul & Burnstein, 1985), and acceptability of a work contract (Golding et al., 1990). However, few other studies (cf. Anderson et al., 1985) have directly examined influence effects when one is presented with a causal alternative as part of the correction. In Experiment 3B, we attempted to extend the results of the previous experiment to additional materials.

Method

Subjects. A total of 40 University of Michigan undergraduates participated in a single session lasting approximately 15 min. Thirty-three of the subjects received course credit for participating, and 7 were paid. Each experimental condition contained both paid and credit subjects. Subjects were run in groups of 2 to 6.

Design. In three conditions, we varied the information presented to subjects. The no-mention control group saw no information about the son as a suspect. The negation and alternative conditions were given information about the son as a suspect, which was later negated in the negating correction group and corrected with an alternative in the alternative correction group.

Materials. The materials were a series of reports describing a police investigation into a theft at a private home. The series consisted of 14 individual messages, each 2 to 4 sentences long (see Appendix D), which were combined into a booklet, with 1 message per page. For the no-mention control group, Message 6 stated that the homeowner's son had been out of town and was unable to watch the house. This condition did not include a correction. For both the correction groups, Message 6 stated that the son was suspected of taking the homeowner's jewelry box to pay off gambling debts, and Message 12 corrected this, stating that the police no longer considered the son as a suspect because he had been out of town when the robbery occurred. Then, in the alternative correction group, Message 13 stated that the police had caught another suspect with the jewelry. In the negating correction group (and the no-mention control group), Message 13 stated that the homeowner would offer a reward for return of the missing jewelry.

A memory test was also prepared, similar to the open-ended questionnaire used in the previous experiments. It included an initial question about the cause of the incident, 10 questions on facts directly presented in the messages, 11 other questions requiring the subjects to make inferences about the event, and 2 final questions assessing whether subjects were aware of the updated information (questions from this measure are shown in Appendix E). All questions appeared in the same order for each subject, with all fact questions appearing before any inference questions to prevent the latter from introducing biases. The 2 questions about noticing the update always appeared after the inference questions.

Procedure. We used the same procedure as in the previous experiments, except that subjects did not fill out a summary sheet or do an intervening task before answering the open-ended questionnaire. The summary sheet was omitted because of time constraints and because results from the fact recall and story recall measures were redundant in the previous experiments.

Results

Scoring. A coder who was uninformed about the experimental conditions scored the responses to the inference questions for a thematic-inference measure and for a global-cause question about responsibility for the theft. Direct references were not coded in this experiment. Two themes were considered: an *inside job* theme, in which the son was implicated, and a *stranger–burglar* theme, in which the son was not thought to be involved with the theft. The inside job theme encompassed references to the son's potential involvement, using a key to enter the house, familiarity with the homeowners or their home, knowing where to find the jewelry, and attempts to conceal the crime from the homeowners. References to a stranger–burglar theme were also coded. These were references that indicated a burglar other than the son, entering the house through a broken window, and not having enough time to steal further items. Table 7 shows sample thematic answers to inference questions from the questionnaire. The analyses reported below exclude idiosyncratic responses and items left blank, which made up 7.7% of the responses to the inference questions.

Inferences. Table 8 shows the mean number of thematic inferences per subject for the inside job and stranger–burglar themes by group. A one-way ANOVA of the number of inferences consistent with the inside job theme showed a significant group difference, $F(1, 37) = 3.41, p < .045, MS_e = 14.88$. Planned comparisons showed a significant difference between the negating correction group and the other two groups, $F(1, 37) = 5.94, p < .02, MS_e = 25.89$. A planned

Table 7
Sample Answers to Inference Questions in Experiment 3B by Theme

Question	Answer	
	Inside job	Stranger–Burglar
What could the Harters have done to better avoid the problem?	Not trusted their son.	Used a safe deposit box.
Why wasn't the TV taken?	Harters would notice it was missing.	It was too heavy; not enough time to take it.
How might the thief have gotten into the house?	Through the door; with a key.	Through the broken basement window.

Table 8
Number of Thematic Inferences Consistent With Misinformation, by Group, From Experiment 3B

Inference	Negation	Control	Alternative
Inside job			
M	3.18	1.70	1.06
SD	3.09	1.75	1.39
Stranger–Burglar			
M	4.09	6.08	7.87
SD	2.51	1.94	1.75

comparison between the no-mention control and the alternative groups was not significant ($F < 1$). Conversely, an ANOVA of the number of inferences consistent with the stranger–burglar theme also showed a significant difference, $F(1, 37) = 11.31, p < .0001, MS_e = 46.96$. Planned comparisons showed significant differences between the negating correction and the no-mention control group, $F(1, 37) = 5.66, p < .023, MS_e = 23.50$, and between this control group and the alternative group, $F(1, 37) = 5.59, p < .023, MS_e = 23.19$. As shown in Table 8, the alternative group endorsed the stranger–burglar theme more than the no-mention control group did, even though they received the same initial misinformation as the negating correction group did.

For the global-cause measure, responses to the direct question about what was responsible for the theft were also scored. No subject in any of the groups answered that the son was responsible, the information initially provided before the correction.

Recall. Fact questions from the questionnaire were scored for accurate content. There were no group differences on this fact recall measure ($F < 1.5$). Subjects recalled a mean of 8.9 facts correctly. For the correction recall measure, a subject was scored as noticing the correction if it was referred to accurately in the open-ended questionnaire. Only two of the subjects failed to report the information that the son was out of town and thus was no longer a suspect.

Discussion

The results replicate the findings in Experiment 3A and suggest that the effects found earlier are not just an artifact of the particular scenario used. As expected, subjects showed a significant influence effect when the correction merely negates previously presented information, but the discredited information has less influence when the correction provides an alternative explanation for events in the scenario. In fact, the group given an alternative after the discredited information appeared similar to control subjects who never heard the original information. This experiment demonstrates that the effect of adding a causal alternative to the correction is not specific to the story used in the earlier studies because the materials in Experiment 3B present a different story context and different features of the account that need explaining. Further, the type of causality differs in Experiment 3B because the scenario deals with an intentional agent (the son) causing events rather than a physical substance (volatile materials) causing events. Thus, the continued influence effect and the mitigation of the effect when a causal alternative is included in the correction are not due to any specific features of the warehouse story used in Experiment 3A.

The two scenarios used in these experiments do share some similarities that may suggest constraints on when the continued influence effect might occur and when corrections with causal alternatives may be effective. First, both stories contain many details that are consistent with the misinformation, even though one can potentially explain those details in another way. One might expect less influence in a story that contains little information other than what was corrected or when the story contains more information that is clearly inconsistent with the misinformation. Second, the events in the stories used here tended to avoid stereotypical default explanations that a subject might readily generate. For example, an auto accident story might tend to be associated with specific schemas in memory (such as drunk driving). If another cause for an accident were discredited, subjects could potentially use this alternative schema already in memory to reconstruct the event and avoid use of misinformation. On the other hand, one might expect more of a continued influence effect if the event is one for which subjects have little structured background information available in memory. Finally, the stories in these experiments do not explain the origin of the discredited information. One generally believes that a communicator is following conversational maxims (Grice, 1975) and thus providing accurate information. If information is corrected later, and one does not decide that the initial communicator was lying, one may still believe that the communicator had a reason for providing the initial information. One may then reinterpret either the correction or the initial information and show continued influence as a result of these reinterpretations. A correction that also explains why a communicator might have provided inaccurate information (e.g., honest mistake, deceit, and so forth) could also reduce the continued influence effect.

General Discussion

The results of the present experiments suggest a reexamination of previous theories about how one comprehends corrections in complex contexts. The results show that misinformation is more likely to influence inferences when it plays a causal role in the account. In Experiment 2, subjects showed many inferences involving volatile materials when the reports presented them as causally sufficient to account for a fire's starting, regardless of whether a correction occurred. However, when the story context made an incidental (noncausal) mention of the same materials, subjects recalled that information but made very few inferences referring to those materials. This provides evidence that a story's causal affordances can determine the degree of influence misinformation may have.

Experiments 3A and 3B demonstrate further effects of causal affordances. Subjects made fewer inferences involving misinformation when the correction also presented a causal alternative to it. This suggests that one may involve misinformation in inferences when it is the only material that affords causal structure to the account, and without it, one could not construct as good a representation of the event. This is consistent with other proposals that subjects use discredited

information when they do not have other "relevant" information available (Anderson et al., 1985; Wegner et al., 1985), but because the mere presence of plausible information does not lead as readily to influence, the current studies specify relevance in terms of the causal role in which the misinformation was presented. These influence effects may be specific to texts in which many of the details support a causal interpretation consistent with the misinformation and the alternative provides plausible causal affordances.

The present experiments also show that two previously proposed accounts of continued influence did not account for the effects found here. First, the results were inconsistent with an availability theory, which predicts that the mere presence of misinformation in the context of the account leads to influence. Subjects in Experiment 2 made few inferences involving the volatile materials information when it was presented incidentally or primed. However, subjects in the noncausal condition recalled the information, suggesting that it was at least available for use. Similarly, subjects in the external reference condition read and generated key words from the volatile materials message other subjects saw. Previous research (Schacter, 1987) has found that exposure to items makes them more likely to appear later on tests that do not require explicit recall of previous information. However, external priming was not sufficient to result in influenced inferences in this comprehension task. Thus, theories that focus on availability of information (e.g., Anderson et al., 1985) do not account for all the influence found here.

Second, the results show that theories based on an "editing failure" assumption (e.g., Ross et al., 1975) also do not account for all the continued influence found in this case. Such theories suggest that one elaborates on misinformation and integrates it into a representation before a correction occurs. Subjects might then have trouble "editing out" these elaborations once the discrediting occurs, and thus could retrieve them from memory later, thereby showing the influence effect. However, another main finding of these experiments is that influence from discredited information can occur in inferences originating after the correction occurs. In Experiments 1A and 1B, subjects who received the correction in a message immediately following the misinformation still showed more influence from it than a control group did and showed the same level of influence as a group with a longer delay between receiving the original message and the instruction to disregard it. However, subjects in the immediate correction group had less opportunity to make precorrection inferences than those in the group with the longer delay, who could potentially make causal and coherence-maintaining inferences before the correction occurred. This work adds to previous research that has suggested that influence occurs as a result of problems in editing prestored inferences (Ross et al., 1975; Wilkes & Leatherbarrow, 1988; Wyer & Budesheim, 1987) by showing that misinformation can also influence postcorrection inferences.

One may also need to consider more complex explanations of how people understand accounts containing corrections. In addition to considering a story's causal affordances, as discussed above, one may want to look at the conversational implications (Grice, 1975) an original presentation of the information has, given that it is corrected later. One tends to assume that communicators have reason to believe any message they provide, before it is shown to be misinformation. One may then have to reconcile one's belief that the communicator was attempting to be accurate with the fact that he or she was actually wrong. This could lead one to reinterpret either the initial information or the correction in a way that leads to continued influence. The materials used in this research do not provide an explanation for how the misinformation originated (e.g., honest mistake, deceit, and so forth); a correction doing so might also be more effective than one that just negates literal content presented earlier. However, our experiments presented the stories in the context of on-line news reports coming in from an investigation scene; consequently, subjects may have more readily accepted updated information that contradicted earlier reports.

Our research extends that of traditional verbal-learning paradigms dealing with directed forgetting (e.g., Bjork & Woodward, 1973), despite seemingly contradictory findings. Our research parallels these directed forgetting paradigms in the emphasis on how people deal with irrelevant information and in subjects' initially assigned task of recall. However, the research here finds continued influence from misinformation, in contrast to the robust forgetting effects found in much of the directed forgetting literature (e.g., Bjork & Woodward, 1973; Horton & Petruk, 1980; MacLeod, 1975). The effects found here may depend on having a coherent, causally related account in which a single or minimal correction has a significant impact on the construal of meaning. Directed forgetting work using sentences (Geiselman, 1974, 1975), as opposed to lists of unrelated words, has also found that subjects forget the forget-cued sentences less effectively when they are thematically related to remember-cued sentences or to each other. The present results suggest further examination of conditions in which people can and cannot seem to deal with irrelevant information successfully.

Our work also differs from that on postevent misinformation and eyewitness testimony (e.g., Loftus, Miller, & Burns, 1978). In that paradigm, subjects witness a complete event, and then receive new, misleading information. In the present research, the information subjects learn initially turns out to be invalid, and they receive new, accurate information during their comprehension of the event. More critically, the postevent misinformation paradigm usually presents misleading information about a peripheral detail in the situation, resulting in subjects being unaware that they have been misled; otherwise, subjects do detect and reject misinformation that blatantly contradicts their perceptions (Loftus, 1979). In contrast, the work in this article shows influence from misinformation that is central to one's understanding of the event, and subjects do openly detect, accept, and accurately report the correction on memory measures.

These results have several implications for theories of comprehension involving corrections. They suggest that asserting information can result in its propagation through later inferences despite even direct, immediate corrections. This in turn suggests that one cannot completely discredit information in memory by merely negating the literal content expressed earlier; instead, an effective correction must go further if it is to result in a successful discrediting. The results also suggest

that the more likely a piece of information is to become involved in causal inferences, the more likely it is to lead to continued influence on inferences after a correction. Thus, one could have more trouble eradicating a central and important piece of information from memory if incorrect, but this would also be the information most important to correct accurately. The correction message should be directly stated and should optimally provide an alternative to the misinformation in order to improve the likelihood of accurate correction in memory. Other potential methods for enhancing the processing of corrections, such as ways of addressing the conversational implications of someone initially providing misinformation, remain to be explored.

In summary, our results support the hypothesis that continued influence from misinformation depends on its propensity for becoming involved in causal inferences and does not depend on simple heightened availability in memory. Further, having a causal alternative associated with the correction of misinformation decreases, but does not completely eliminate, its influence. The results also show that subjects' inferences about an event can show influence from misinformation, even when they encounter an immediate, direct instruction to disregard it. Finally, the results suggest that such inferences can originate after the correction occurs in which subjects use the misinformation at a later test time to generate new inferences. This contrasts with previous work that has posited continued influence due to failure to edit prestored inferences (Anderson et al., 1980; Wyer & Budesheim, 1987) and suggests that two different processes may result in continued influence from misinformation.

References

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 & Cognition, 15,* 247–255.

Bjork, R. A., & Woodward, A. E. (1973). Directed forgetting of individual words in free recall. *Journal of Experimental Psychology, 99,* 22–27.

Carretta, T. R., & Moreland, R. L. (1983). The direct and indirect effects of inadmissible evidence. *Journal of Applied Social Psychology, 13,* 291–309.

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.

Geiselman, R. E. (1975). Semantic positive forgetting: Another cocktail party problem. *Journal of Verbal Learning & Verbal Behavior, 14*(1), 73–81.

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.

Graesser, A. C. (1981). *Prose comprehension beyond the word.* New York: Springer-Verlag.

Graesser, A. C., & Clark, L. F. (1985). *Structures and procedures of implicit knowledge.* Norwood, NJ: Ablex.

Grice, H. P. (1975). Logic and conversation. In P. Cole, & J. L. Morgan (Eds.), *Syntax and semantics, 3: Speech acts* (pp. 41–58) San Diego, CA: Academic Press.

Hastie, R., & Park, B. (1986). The relationship between memory and judgment depends on whether the judgment task is memory-based or on-line. *Psychological Review, 93,* 258–268.

Haviland, S. E., & Clark, H. H. (1974). What's new? Acquiring new information as a process in comprehension. *Journal of Verbal Learning and Verbal Behavior, 13,* 512–521.

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

Horton, K. D., & Petruk, R. (1980). Set differentiation and depth of processing in the directed forgetting paradigm. *Journal of Experimental Psychology: Human Learning and Memory, 6*(5), 599–610.

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.

Kintsch, W. (1974). *The representation of meaning in memory.* Hillsdale, NJ: Erlbaum.

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

Loftus, E. F. (1979). Reactions to blatantly contradictory information. *Memory & 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.

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

Massad, C. M., Hubbard, M., & Newtson, D. (1979). Selective perception of events. *Journal of Experimental Social Psychology, 15,* 513–532.

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

O'Brien, E. J., & Albrecht, J. E. (1991). The role of context in accessing antecedents in text. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 94–102.

O'Brien, E. J., & Myers, J. L. (1987). The role of causal connections in the retrieval of text. *Memory & Cognition, 15,* 419–427.

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.

Schacter, D. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13,* 501–518.

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

Spiro, R. J. (1980). Accommodative reconstruction in prose recall. *Journal of Verbal Learning and Verbal Behavior, 19,* 84–95.

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 Graesser, A. C., & Bower, G. (Eds.), *The psychology of learning and motivation: Inferences and text comprehension* (Vol. 25, pp. 175–196). San Diego, CA: Academic Press.

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

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

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 and Social Psychology, 48,* 533–549.

Appendix A

Materials Adapted From Wilkes and Leatherbarrow (1988)

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 (Mention and Discrediting Groups)[A1]

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

Message 5 (Control)[A1]

4:30 a.m. Message received from Police Investigator Lucas saying that they have reports 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 (Control)[A1]

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 11 (Discrediting Groups)[A1]

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 oil paint and gas cylinders had actually been empty before the fire.

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 photocopying 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.

Note. The materials are adapted from "Editing Episodic Memory Following the Identification of Error" by A. L. Wilkes and M. Leatherbarrow, 1988, *Quarterly Journal of Experimental Psychology, 40A,* pp. 381–383. Copyright 1988 by The Experimental Psychology Society. Adapted by permission

[A1] These are the messages that changed in the different conditions.

Appendix B

Comprehension Questions

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

1. Why did the fire spread so quickly?
2. For what reason might an insurance claim be refused?
3. What was the possible cause of the toxic fumes?
4. What was the relevance of the closet?
5. What aspect of the fire might the police want to continue investigating?
6. Why do you think the fire was particularly intense?
7. What is the most likely cause of the fire that the workers successfully put out earlier?
8. What could have caused the explosions?
9. Where was the probable location of the explosions?
10. Is there any evidence of careless management?

Manipulation Check Questions

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

Appendix C

New Messages Added to Story

New Message Added in Experiment 2

Message 3 (noncausal reference group)[C1] The alarm was raised by the manager of a store located across the street, which sells cans of oil paint and pressurized gas cylinders. He had smelled smoke and become concerned.

New Messages Added in Experiment 3

Message 5. The fire marshal had recorded several fire code violations on the premises at a surprise inspection two months earlier.

Message 7. The display room was reported to contain display cases, catalogs, and the sales staffs' desks. It was only staffed from 11 a.m. to 2 p.m., due to diminishing sales.

Message 14 (alternative and arson control groups).[C1] 11:08 a.m. Firefighters have found 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. The owner denies any knowledge of these materials.[2]

[C1]These are the messages that changed in the different conditions.

Appendix D

Jewelry Theft Story

Message 1

3:00 p.m., May 2nd. Police respond to a call made from a home on Acorn St., in a middle-class residential neighborhood.

Message 2

The homeowner, Ms. Harter, reports that her jewelry box is missing. Contents are reported to include gold chains, gold and silver earrings, rings, and pendants of precious stones.

Message 3

She discovered that the box was missing when she returned from a vacation and wanted to put a new necklace she'd bought in it. It had been stored in a locked drawer in her bedroom dresser.

Message 4

She swears she had checked the box before leaving on vacation, and everything was in order. A tall tree arches near the bedroom window, but police have found no evidence of tampering with the window.

Message 5

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.

(Appendixes continue on next page)

Message 6 (Control)[D1]

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.

Message 6 (Correction)[D1]

Police suspect that Evan may have taken the box from the house to help pay off large gambling debts.

Message 7

The neighborhood has been hit with a number of thefts recently. There are no arrests or leads in these cases so far.

Message 8

The Harters' next-door neighbor reports that she noticed a light on in the 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 9

A search for footprints and tire tracks has turned up inconclusive, due to a recent rainstorm. In the course of the investigation, an officer noted a broken latch on a basement window.

Message 10

Police are still attempting to determine whether other valuables are missing from the house. The television and a home computer had not been disturbed, however.

Message 11

The Harters have contacted their insurance company about the loss. The last appraisal showed the box's contents to be worth several thousand dollars.

Message 12 (Control)[D1]

A second message from the police investigators about the incident. It stated that several officers had checked the main pawn shops in town, and confirmed that no jewelry matching the Harters' descriptions has surfaced yet.

Message 12 (Correction)[D1]

A second message from the police investigators about the incident. It stated that the Harters' son is no longer a suspect, because several independent sources confirm that he had been out of town on business during the Harters' vacation.

Message 13 (Control)[D1]

Ms. Harter is considering offering a reward for return of several of the pieces, because they have great sentimental value for her. There would be no questions asked.

Message 13 (Alternative)[D1]

Police have now arrested excon Dan Fowler, who had tried to pawn some of the jewelry. His girlfriend, Ana, had told him about Ms. Harters' jewels, which she had heard about while working as a new maid in the neighbor's house.

Message 14

Detectives will look for similarities between this case and the other thefts reported in the neighborhood recently.[3]

[D1]These are the messages that changed in the different conditions.

Appendix E

Comprehension Questions

Cause Question

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

Fact 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 Ms. Harter consider offering a reward?
6. What did the Harters' next-door neighbor notice?
7. What kinds of jewelry did the box contain?
8. What arrangements did the Harters make for checking up on the house?
9. When did Ms. Harter discover that the jewelry box was missing?
10. What did police notice about the bedroom window?
11. When did the neighbor's dog suddenly start barking?

Inference Questions

12. Why might the neighbor's dog have been barking?
13. Whose car might the neighbor have noticed, parked in the alley?
14. Why might the son feel bad about the incident?
15. What could the Harters have done to better avoid this problem?
16. Who, if anyone, should be questioned more thoroughly by the police?
17. Why wasn't the television taken?
18. How might the thief have gotten into the house?
19. Why might the Harters be angry with their son?
20. What might be responsible for the other thefts in the neighborhood recently?
21. What steps should the police take next?
22. Where was Evan Harter on the evening of April 28th?

Manipulation Check Questions

23. What did the police investigators 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?

Received July 27, 1992
Revision received December 23, 1993
Accepted December 30, 1993 ■