

A theoretical review of the misinformation effect: Predictions from an activation-based memory model

MICHAEL S. AYERS and LYNNE M. REDER
Carnegie Mellon University, Pittsburgh, Pennsylvania

This article reviews the major empirical results and theoretical issues from over 20 years of research on people's acceptance of false information about recently experienced events (see, e.g., Loftus, 1975). Several theoretical perspectives are assessed in terms of their ability to account for the various and sometimes conflicting results in the literature. Theoretical perspectives reviewed include the trace alteration hypothesis, the blocking hypothesis, the task demands/strategic effects hypothesis, source monitoring, and an activation-based semantic memory account. On the basis of its ability to account for the reviewed data and other cognitive phenomena, an activation-based semantic network model of memory is suggested for understanding the data and planning future research in the area.

In her seminal work on the topic of the malleability of memory, Loftus opened an area of study that has persisted for over 20 years (Loftus, 1975; Loftus & Palmer, 1974). The basic result was that exposure to misleading questions about an experience often resulted in an apparently permanent loss from memory of the original details, which were replaced by the false information. This phenomenon has come to be known as the *misinformation effect*. Loftus's interesting and somewhat surprising result sparked a flurry of papers by many authors, some supporting the original claim and others explaining aspects of the data by means other than memory impairment—such as strategic effects, guessing biases, or task demands. The primary question has been whether and, if so, to what extent memories for events can be changed by subsequent exposure to misinformation about those events. The answer to this question has important implications, both for memory theorists and for the general public.

Perhaps the most salient social implication concerns the veridicality of eyewitness testimony in criminal trials. Eyewitness testimony is generally accepted by jurors as one of the most compelling types of evidence. If, under some circumstances, witnesses' memories may differ from the actual events, jurors should be made aware of this. Several good discussions on this subject and other related social issues already exist (see, e.g., Ceci & Bruck, 1993; Lindsay & Read, 1994; Loftus, 1991; Wells & Turtle, 1987).

To date, dozens of papers investigating different aspects of the misinformation effect have been published. Despite this wealth of data, none of the proposed theoretical explanations adequately accounts for the bulk of the findings. The primary aims of this paper are to review the kinds of explanation that have been given for the misinformation effect and to suggest a single theoretical framework within which to account for the various and apparently conflicting results in the literature. In keeping with this focus on discrepancies among accounts, the paper is organized roughly chronologically—that is, in the order in which the different explanations have arisen over the past two decades. The discussion is limited primarily to the issues, problems, or inconsistencies that are particularly troubling for one or more of the theoretical perspectives. We begin by reviewing the original paradigm and results.

THE ORIGINAL PARADIGM

For purposes of comparison, it is important to specify the original procedure employed by Loftus and her colleagues. In the early experiments on the misinformation effect (see, e.g., Loftus, 1975, 1979; Loftus, Miller, & Burns, 1978), the subjects first watched a video tape or slide show of an event, such as a traffic accident. As an example, suppose that one of the slides depicted a car at a stop sign. After the slide sequence, the subjects were asked a series of questions about the events in the slides. Embedded in one of these questions was the misleading presupposition that the car was stopped at a *yield* sign, even though the slides had shown a stop sign (e.g., "Did another car pass the red Datsun when it was stopped at the yield sign?").¹ When later tested on the information seen in the slides, the subjects who received the misleading question were much more likely to report having seen a yield sign than were the subjects whose intervening questions contained correct information (*stop sign*) or neutral information (*intersection*).

Preparation of this manuscript was supported by Grant N00014-95-1-0223 from the Office of Naval Research, Grant 1R01 MH52808-01 from the National Institutes of Mental Health, and AFOSR Grant F49620-97-1-0054. The authors wish to thank Brad Best for his comments on an earlier version of the manuscript and Alison Clark for her comments on all versions of the manuscript. M.S.A. is now at the School of Education, University of Pittsburgh. Correspondence concerning this article should be addressed to L. M. Reder, Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213 (e-mail: reder@cmu.edu).

This increased likelihood to report having seen something in one part of an experiment that was only suggested (implied as a presupposition) in a different part of the experiment is known as the misinformation effect. Most experiments investigating the misinformation effect have used some variant of this original paradigm.

Loftus and Palmer (1974) found that the wording of a suggestive question could influence the likelihood of people's reporting that they saw things that had not occurred. Subjects who read a question suggesting that the cars in the film *smashed* into each other were more likely to report seeing broken glass than were subjects whose questions suggested the cars *collided*, *bumped*, *contacted*, or *hit*, even though no broken glass appeared in the film. These data suggested that people's memory reports were not simply mental replays of the events experienced but instead were a reconstruction that was susceptible to intervening information and errors of inference.

Several kinds of explanations have been offered for these and subsequent results. Some of these implicate memory impairment, suggesting either alteration of the original memory traces by the intervening information or competition among multiple traces. Other explanations focus on factors other than memory impairment—such as strategic effects and task demands. Still others focus on the vulnerability of source information, which are subject to decay or interference. The remainder of this article discusses each of these classes of explanations, noting their strengths and weaknesses for a successful account of the relevant data. At the same time, we compare these accounts with our framework, which embodies a few mechanistic assumptions. We believe this framework incorporates some of the strengths of the various theories and provides an account of some otherwise inconsistent findings with a general activation-based memory model.

To foreshadow our discussion, it is well established that each of the perspectives to be discussed has contributed substantially to our understanding of people's susceptibility to misinformation. We use the term *misinformation effect* for differences in performance between misled and control conditions, regardless of the mechanism responsible for these differences. The misinformation effect sometimes reflects true memory impairment, sometimes stems from various nonmemory phenomena—such as task demands, strategic effects, and response bias—and sometimes is a combination of these effects. To distinguish between these classes of factors or mechanisms, we use the term *memory impairment* to refer to the memory-based component of the effects discussed; for the nonmemory-based contributions to the effects discussed, we refer specifically to the most likely factor (e.g., task demands, strategic effects, or response bias).

The framework offered in this article accounts primarily for the effects of memory impairment. The discussion of the data in the context of a model helps to illustrate the interactions among memory-based and nonmemory-based causes for the misinformation effect. In our view, the memory impairment effects frequently found in misinforma-

tion effect studies are essentially the same effects that have been studied in verbal-learning paradigms for decades (see, e.g., Anderson, 1974a; Barnes & Underwood, 1959; McGeoch, 1942; Melton & Irwin, 1940; Postman & Underwood, 1973; Roediger & Schmidt, 1980; Smith, 1971; Underwood, 1965). Clearly, ours is not the first attempt to account for interference data on the basis of competition of associations. However, our approach to understanding the misinformation effect literature is unique in that it is based on a mechanistic model.

In brief, we have three main goals in this review: (1) to provide a mechanistic account of the memory-based effects in the misinformation effect literature, (2) to provide a framework for distinguishing memory-based from nonmemory-based factors that influence people's performance in these tasks, and (3) to draw relevant ties between this literature and the verbal-learning literature in order to provide a more general perspective on this body of research. We begin by introducing a specific activation-based framework, to which we will later refer in our discussion of the critical issues.

THE ACTIVATION-BASED FRAMEWORK

Our theoretical account assumes an activation-based semantic network model of memory (see, e.g., Anderson, 1976, 1983, 1993; Anderson & Bower, 1973; Bower, 1996; Collins & Loftus, 1975; Kamas & Reder, 1994; Reder & Gordon, 1997; Reder & Schunn, 1996; Schunn, Reder, Nhouyvanisvong, Richards, & Stroffolino, 1997). Models that assume this type of representation have been used to account for a variety of memory phenomena—including associative priming (Ratcliff & McKoon, 1981), the fan effect (Anderson, 1974b), feeling-of-knowing (Reder & Schunn, 1996; Schunn et al., 1997), the shape of learning functions (McClelland & Rumelhart, 1986), word frequency effects in recognition (Reder et al., 1998; Reder et al., 1997), among others—while resorting to relatively few ad hoc assumptions. The various models of this class vary in their terminology and specific assumptions; so, for clarity, our discussion focuses on one specific instantiation, the source of activation confusion (SAC) model of memory (Reder & Schunn, 1996; Schunn et al., 1997).

Representation of concepts and associations. Figure 1 is a schematic representation of the way in which SAC assumes that information is stored in memory. In the figure, concepts are represented by the ovals, and associations between concepts are represented by the lines connecting the ovals. A concept's representation in memory gets stronger or weaker as a function of how frequently and how recently it has been encountered. In Figure 1, we represent the various concepts' current strength (or resting level of activation) by the relative thickness of the ovals. Relations among concepts also vary in strength as a function of their frequency and recency of co-occurrence. Strength of relations is represented in Figure 1 by the thickness of the lines connecting the concepts.




Figure 1. Partial schematic representation of slide show and narrative used in some misinformation effect experiments. This is an illustration of part of the memory structure used by activation-based models. Nodes (ovals) represent semantic concepts and links (lines) represent associations between concepts.

The long-term growth and decay of concepts' and associations' strength are described by a power function; however, there is also a rapid increase or decrease in current activation, which is described by an exponential function (see Reder et al., 1997; Reder & Schunn, 1996; Schunn et al., 1997, for details of the model and specific equations). Availability of a concept is a function of its current activation. Not only does prior exposure affect this value, but concepts also receive activation from associated concepts when the activation level of these potential sending concepts exceeds a critical threshold. The amount of activation a concept receives from an associated sending or source concept is a function of the strength of that connection relative to the strength of competing associations from the sending node (the fan off the source node). In Figure 1, the concept *yield sign* is stronger than the concept *stop sign* because it has been seen more recently.

Representation of context and encoding source. Another issue of relevance is the representation of context. In principle, context could be represented as a vast web of concepts and associations that includes such things as details about the room in which the experiment takes place, particularities of the experimenter, the subject's emotional state, and so forth. For simplicity, context is represented in Figure 1 as a few schematic bundles: one corresponding to the features associated with the oral presentation of the narrative (if the item was given in the narrative); another bundle representing the features associated with the visual slide presentation (if the item was shown in the slides); and a third bundle schematically representing all the remaining features of the experimental context (i.e., those associations not unique to either the slide sequence or the narrative).

We assume that, as concepts are encountered and stored in memory, associations to contextual features are automatically created; however, the number of contextual associations created may vary with task demands, attention, and the richness of the encoding context. Whether these contextual features are accessed by (or even accessible to) the subject at the time of test is determined by a number of variables, some of which are represented directly in the model—for example, the number of other concepts to which the accessed item is associated, the relative strength of those associations, and the number of other associations to the context—and others of which are not represented directly—for example, the subject's perceived task demands and motivation to be accurate. This issue is discussed in more detail in the next section.

Source of activation confusion. The model is called *source of activation confusion* (SAC) because we believe people's actions are affected by *whether* a concept is active, even though they are often unaware of *why* a concept is active. In terms of the schematic representation, people have access to the concepts themselves (the ovals), but have no direct access to the links along which activation travels (the lines). People can only try to infer the source responsible for a concept's activation.² A concept might be active for a variety of reasons—for example, because it received activation from an associated concept, because it was subliminally flashed, or because it was just studied. Under some conditions, the source of a concept's activation is misattributed, which results in a memory error (Kelley & Jacoby, 1996).

Figure 2 provides a slightly more detailed schematic representation of the critical elements in the original misinformation paradigm. Note that, associated with the crit-

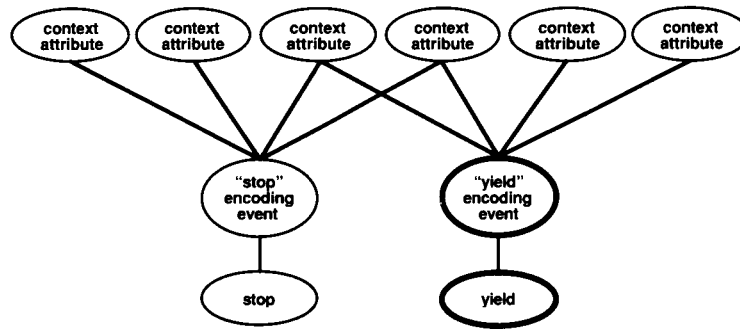


Figure 2. Slightly more detailed schematic representation of the critical aspects of Figure 1 in determining *remember* responses.

ical concepts (stop sign and yield sign), there are nodes that represent the encoding events for those items. Although not depicted in the figure, the model assumes a similar encoding event for each concept encountered. Associated with each encoding event is a set of concepts that represent contextual features. For simplicity, in Figure 1 these features are combined and represented as a single contextual association. For many of the early misinformation effect phenomena, it is not necessary to discuss the representation at the level of greater detail; the level presented in Figure 1 is sufficient. More recently, however, much emphasis has been placed on determining whether subjects actually *remember* the suggested details. In order to account for these data, a more detailed schematic representation is necessary, although the formalization of the model would be the same in either case. We should note that this assumption about the representation of contextual information was not made specifically to account for the data reviewed here; it is a fundamental assumption of the model and has been used to formally model other experimental data (see Reder et al., 1998; Reder et al., 1997).

Summary of the model. To review, the major assumptions of the model are: (1) a concept's strength decays over time (since its last exposure) in accordance with a negative power function; (2) activation spreads from a given concept to all associated concepts in proportion to their strength of association, which is determined by their recency and frequency of co-occurrence; (3) context is represented as a collection of concepts representing distinct features of the encoding situation; (4) people have conscious access to concepts themselves, but not to the associations among them; and (5) people sometimes make errors when attempting to infer the source (cause) of a concept's current level of activation.

THE OVERWRITE/TRACE ALTERATION ACCOUNT

The original interpretation of the basic result was that the memory trace was *altered* or *overwritten* by the subsequent suggested information (see, e.g., Loftus, 1975, 1979; Loftus et al., 1978). This account assumes that, once

the misleading information is encoded, the original trace is overwritten and no longer exists in its original form.³ When subjects are later asked about the original information, they instead access this updated trace containing the misinformation.

According to the trace alteration account, memory traces are created as the subject views the slides. One of these traces represents the originally seen stop sign. Later, the subjects are asked a misleading question that refers to a yield sign. The trace alteration account posits that, when the yield sign is mentioned, the memory trace for the stop sign is sometimes overwritten by this new yield sign trace. Thus, on the final test, some of the subjects who received the misinformation no longer have the stop sign in memory because it has been replaced by the yield sign.

This explanation differs from our activation-based account in that it assumes a single trace, whereas, in our view, an additional trace is added to the concept *sign* that is attached to the representation of the scenario depicted in the slide sequence. Now that there are two traces attached to *sign*, one for the original and one for the misleading information, activation is shared by the traces, so either could be given as a response—the response given is the one with greater activation at the time of the test. This contrast between single-trace and multiple-trace accounts is important and reappears as an issue in the discussion of other findings. The field appears to be approaching a consensus that memory impairment effects in the misinformation paradigm are best explained by some form of multiple-trace account, although the specific form is still an issue of debate (see Chandler, 1991; Chandler & Gargano, 1995; Windschitl, 1996).

In many ways, this issue of single- versus multiple-trace accounts of the misinformation effect is similar to the debate between theories of forgetting based on interference as opposed to those based on unlearning—the former explanation being based on response competition and the latter assuming impairment of an earlier memory trace. McGeoch (1932, 1942) advocated a theory of forgetting based on competition of multiple responses associated with a common stimulus to account for people's impaired memory performance on a list when there had been an in-

tervening list between study and test, as compared to the same temporal delay without an intervening list. Melton and Irwin (1940) argued that, although response competition was a contributing factor, unlearning of the original items during the study of subsequent items must also occur in order to explain the fact that most errors were errors of omission rather than intrusions from an intervening study list (see Crowder, 1976, for a more complete review).

THE STRATEGIC EFFECTS ACCOUNT

The trace alteration hypothesis interprets the misinformation effect in terms of a change to the structure of memory per se (the original trace is replaced). In contrast to this account, the strategic effects account explains the misinformation effect in terms of factors other than memory impairment (see, e.g., McCloskey & Zaragoza, 1985a, 1985b; Zaragoza, McCloskey, & Jamis, 1987). Such factors include demands of the type of final test used and subjects' use of process-of-elimination strategies when guessing.

McCloskey and Zaragoza (1985a) reasoned that the apparent memory impairment found in the first misinformation effect studies might not be due to real changes in the subjects' memory for the critical information but instead might result from task demands. They suggested that subjects sometimes simply forget the critical information that was presented, regardless of whether or not misinformation was also presented. The original testing procedure does not allow the experimenter to determine whether the misinformation effects are due to interference from the introduction of the new information or whether there is a strategic bias to select the misinformation in the absence of other memories. That is, if the critical information is no longer recognized, but there is an alternative that is recognized as having been presented in the narrative, it might seem most sensible to the subject to select the item that was in the narrative. Of course, unbeknownst to the subject, such a bias would systematically produce the wrong response.

A second possible explanation that McCloskey and Zaragoza (1985a) offered for the misinformation effect is that the subjects may remember both the original and the misleading information but choose to report the item from the misleading narrative in order to produce the results the subject believes are wanted by the experimenter. In other words, the results from studies using the original test may be due, at least in part, to task demands rather than to true memory impairment.

The Modified Recognition Test

In order to control for possible effects of bias or task demands, McCloskey and Zaragoza (1985a) developed a modified testing procedure. Like Loftus's original procedure, the modified procedure consists of a series of slides, followed by a written narrative containing misinformation about some details depicted in the slides. The modification involves changing the alternatives on the final memory test. McCloskey and Zaragoza (1985a) developed a set

of stimuli that would be used in many studies investigating the misinformation effect. The materials consisted of a slide show depicting an office robbery. As in the procedure developed by Loftus and her colleagues, the subjects in McCloskey and Zaragoza's (1985a) experiments first viewed the slide show and then received misinformation, either imbedded in questions or in a narrative describing the series of events from the slides. One of the critical slides showed a man placing an object under a *hammer* in a tool box. The misinforming narrative or questions suggested that the item was placed under the *wrench*, even though no wrench appeared in the pictures. Except for the set of materials used, this part of the modified procedure is identical to Loftus et al.'s original procedure.

The critical difference between the original procedure and McCloskey and Zaragoza's (1985a) modified procedure is the type of final memory test given. Instead of asking the subjects to discriminate the originally seen information from the misleading information (e.g., *hammer* vs. *wrench*), the modified test excludes the misleading information as an alternative. Instead, the two choices are the original item and a new item the subject has not seen at all (e.g., *hammer* and *screwdriver*). McCloskey and Zaragoza (1985a) argued that, if exposure to misinformation impairs subjects' memory for the original information, subjects who received misinformation still should be less likely than are control subjects to correctly identify the item originally seen in the slides. Studies with this modified procedure have often failed to find a difference in performance between misled and control conditions (Belli, 1993; Bowman & Zaragoza, 1989; Ceci, Ross, & Toglia, 1987; Chandler, 1989, 1991; Loftus, Donders, Hoffman, & Schooler, 1989; McCloskey & Zaragoza, 1985a). However, it also should be noted that several studies have found effects of memory impairment with this procedure (Belli, Windschitl, McCarthy, & Winfrey, 1992; Ceci et al., 1987; Chandler, 1989, 1991; Windschitl, 1996). Findings of null effects are consistent with the strategic effects account, which predicts that the misinformation effect found with the original testing procedure should be eliminated when the test does not include the misinforming item as an alternative.

Further support for the strategic effects account comes from experiments using recall rather than recognition for the final test. As in the modified procedure, the purpose of this procedure was to assess impairment of the original item, independent of the state of memory for the misleading item. Zaragoza et al. (1987) used a procedure in which the subjects viewed a slide show of an office robbery, then read either a neutral or misinforming narrative, and finally were given a cued recall test. The misleading narrative differed from that of typical misinformation experiments in that the suggested item came from a category different from that of the original item. For example, if the critical item from the slides was a soft drink can, the misleading narrative might suggest that it was a Planters peanuts can. The question on the final test would then specify the category to which the original item belonged—for example, "The keys to the desk drawer were next to a soft drink can.

What *brand of soft drink* was it?" Thus, in a manner similar to that of the modified procedure, this cued recall procedure prevented the subjects from reporting the misleading information. Zaragoza et al. reasoned that, if misinformation truly impairs memory for the original trace, the subjects should be less likely to retrieve the correct item after receiving misinformation than after receiving neutral information. At both low (around 30%) and moderate (around 60%) levels of recall performance, Zaragoza et al. found no evidence for memory impairment with cued recall, but did find impairment with the original Loftus et al. (1978) two-alternative forced-choice procedure. On the basis of these results, Zaragoza et al. concluded that the difference between misled and control subjects in the original procedure was due to factors other than memory impairment.

According to the strategic effects account, no misinformation effect is found with the cued recall test because the biases present in the original recognition test are eliminated. The guessing bias discussed in previous sections of this paper does not apply in the retrieval paradigm because the misleading item is disallowed by the form of the question and subjects must generate all their responses. These results clearly indicate that strategic effects and biases account for a substantial part of the misinformation effect under some circumstances.

McCloskey and Zaragoza (1985a) correctly reasoned that, if there were no memory for the original information, it would be plausible and sensible to select the misleading alternative; however, our model predicts a small but real effect even when the misleading item is disallowed on the final test. This prediction stems from the assumption that, even though the misleading term is not re-presented as an alternative at test, it still exists in memory and competes with the correct response. As is discussed below, there is evidence for the small effect predicted by our model; furthermore, there are reasons to believe the modified test may be insensitive to impairment in some situations.

Beyond Strategic Effects: True Memory Impairment

Payne, Tolia, and Anastasi (1994) conducted a meta-analysis of 44 published experiments that used modified recognition tests. Individually, only 14 of the 44 experiments yielded a significant misinformation effect. Overall, however, 30 of the 40 nontied cases showed a misinformation effect, and the average effect size of all 44 cases was 4.1%—both being significant effects. This finding clearly supports SAC's prediction of a small memory impairment effect in studies using the modified recognition paradigm.

Response time data. Although it is clear that subjects are sometimes influenced by strategies and biases stemming from the structure of the task, the findings of Loftus et al. (1989) suggest that this explanation is insufficient to explain the misinformation effect. Loftus et al. (1989) reported response times (RTs) that are conditional on whether subjects were correct or misled. For the classic paradigm, control subjects responded more quickly when they were

correct than when they were incorrect, but the reverse was the case for misled subjects. Misled subjects were faster to respond when they chose the misinformation than when they chose the correct alternative.

According to Loftus et al. (1989), the strategic effects account has difficulty accounting for these RT data. They argued that, were it the case that misled subjects used some process of deliberation for choosing between the two items on the recognition test, they would be slower than control subjects, because this deliberation should take time. The data do not bear this out; if anything, misled subjects' error RTs were slightly faster than were the RTs for control subjects' correct responses (Loftus et al., 1989, Figure 1). The RT data appear to indicate that subjects do not often engage in deliberative processes between the alternatives at the time of the test. In other words, the strategic effects account may account for some of the difference between the misled and control conditions, but there is certainly more to the story than can be explained by task demands alone. It is now clear that both strategic effects and memory impairment are involved in the misinformation effect. The question remains, however, as to what variables determine the extent to which each plays a role.

What determines whether an effect is detectable with the modified test? The modified recognition test (McCloskey & Zaragoza, 1985a) sometimes yields impairment effects and sometimes does not. Two issues have been raised with regard to the failure of some studies with the modified testing procedure to show a misinformation effect. The first issue concerns the ability of the modified test to exhibit memory impairment at all. Under conditions of low overall memory performance or only modest impairment, the modified test may be insensitive. The difference between the experimental and control conditions in the modified test is expected to be only half the true size of the impairment effect, because subjects who do not remember the original information still have a 50% chance of choosing the correct information by guessing (Chandler, 1989; Loftus, Schooler, & Wagenaar, 1985). The second issue concerns the sensitivity of the test in the detection of the kinds of memory impairment that depend on the presence of the misinformation at test (Belli et al., 1992). The modified test is only capable of detecting memory impairment that does not require the additional strengthening of the misinformed information that occurs when the misinformation is given as an option at test. Belli et al. (1992) provide a good discussion of both of these issues.

Effects of overall memory performance. Chandler (1989) presented subjects with a series of slides of scenes (e.g., landscapes, lilies floating in a pond). The subjects were shown one slide, then a similar one, and finally were given a modified recognition test consisting of the original slide and a third, new slide that was similar to the first two shown. Chandler found memory impairment with the modified testing procedure in Experiments 1 and 2, in which memory performance was relatively high, but not in Experiments 3 and 4, in which memory performance was near chance. The two critical implications of this se-

ries of experiments are: (1) that genuine memory impairment can be detected with the modified test and (2) that the null results of previous studies with the modified test may have been due to floor effects. The Payne et al. (1994) meta-analysis further supports the latter conclusion. Examining 44 published experiments that employed the modified test, Payne et al. found a significant correlation ($r = .54$) between overall recognition performance and effect size. The strategic effects account predicts no misinformation effect in the modified test, regardless of overall memory performance. The presence of an effect with this different set of materials is best explained as being a result of memory impairment.

Effects of retention interval. Belli et al. (1992) presented further evidence for memory impairment effects. In a series of experiments with the modified test procedure, they found a misinformation effect at long (5–7 days) but not short (15 min) retention intervals. This study used centrally rather than peripherally presented critical items—that is, the critical items were clearly the main subject of the pictures rather than seemingly less important details—in order to boost overall memory performance. Because the test was identical at short and long retention intervals, nonmemory-based explanations have difficulty accounting for the change in performance. In neither case was the misinformation reintroduced at the time of the test, so the strategic effects account predicts no effect at either retention interval. There is no a priori reason to assume that strategies and biases change systematically with delay.

Payne et al. (1994) noted, however, that, although retention interval appeared to be an important factor in determining whether a misinformation effect would be found with the modified test, the few experiments examining this issue (Belli et al., 1992) had necessarily confounded retention interval and encoding conditions in order to avoid floor effects for overall recognition in the delayed-test conditions. A more recent series of experiments by Windschitl (1996) used a modified recognition procedure and rectified the problem of retention interval being confounded with encoding conditions. In these experiments, subjects first were shown a series of faces, then were shown an interpolated series of similar faces, and finally were given a modified recognition test on which they had to choose between the original face and a completely new but similar face. For control items, no interpolated face was presented. In all phases of the experiment, labels were presented below the pictures so that, on the final test, the subjects could more easily determine which faces they were supposed to be comparing (e.g., there was a *bride* picture in each phase of the experiment, which was clearly distinguishable from all the other pictures in that phase).

Across his Experiments 1–3, Windschitl (1996) used retention intervals ranging from 10 min to 2 weeks. He found impairment effects with retention intervals of 10 and 45 min, but not at delays of 48 h or 1 week.⁴ Furthermore, Windschitl's data suggest that the lack of an effect at the 48-h delay could not be explained by a functional floor effect, because overall recognition performance was

significantly poorer at the 1-week delay than after 48 h. Thus, any theory that, it was hoped, could explain the disappearance of an effect with delay could not do so simply by positing that subjects no longer had sufficient memory of the original pictures to detect an impairment effect; a tenable theory must predict that actual impairment decreases over time, independent of simple forgetting. Windschitl conducted power analyses, which further supported this interpretation.

One aspect of Windschitl's (1996) data is particularly challenging for several accounts. In Experiment 2, as expected, the subjects' overall recognition performance for control items worsened as the retention interval increased from 45 min to 48 h (74.8% vs. 69.1%, respectively). Somewhat surprisingly, however, *absolute* recognition performance for experimental items (those for which an interpolated face was shown) *increased* from 63.2% after 45 min to 68.6% at the longer delay.

Windschitl (1996) suggested that this rebound is particularly troubling for a trace alteration account because it would require that the original memory trace, which was altered by the interpolated image, must somehow recover partially from this alteration. To our knowledge, no such mechanism has been proposed. The very presence of an effect with the modified procedure suggests that there is something more than biases or strategic effects associated with viewing the interpolated face. Windschitl's data are most easily explained in terms of retrieval-based memory impairment.

The SAC model accounts for the effect of improved absolute recognition performance at longer delays as follows. Recall that one of the assumptions of the model is that activation of a concept decays with time in accordance with a negative power function: rapidly at first and then more slowly as time progresses. Thus, at a short retention interval, the representation of the interpolated item is strong relative to the originally seen item because it has been seen more recently. At test, even though the interpolated item is not presented as an alternative, it still receives activation from the queried concept (e.g., the *bride* concept)—activation that, in the control condition, goes only to the originally seen item. Thus, in comparison with the control condition, the original item in the experimental condition receives less activation and is less often correctly recognized. At the longer retention interval, the traces for the interpolated and original faces are at roughly equivalent points on their forgetting curves. In other words, the interfering effect of the interpolated item is proportionally muted because it no longer has a recency advantage over the original information. In comparison with the shorter delay, a smaller proportion of the activation from the schema goes to the interpolated item. Consequently, a greater proportion of the activation from the schema goes to the original item than at the shorter delay, which results in an apparent "spontaneous recovery" of the original information.

This explanation for the spontaneous recovery of original information is reminiscent of Crowder's (1976) interpretation of Miller and Stevenson's (1936) study of habit

learning and extinction in rats with an AB-AD design. Crowder argued that, at short delays, impairment effects on the B response to the A stimulus can be explained by interference or competition with the interpolated D response to A, which was learned on an intervening list. The reduction in impairment effects over time is a function of the new association's relatively steep forgetting curve, compared with that of the original association, which is already near asymptote. Although Crowder was explaining cued recall data, the underlying mechanism for the change in effect sizes with increasing delay is the same as that for the modified recognition test in Windschitl's (1996) experiments just discussed. Recently, Wheeler (1995) also demonstrated spontaneous recovery and made arguments similar to our own.

Using an AB-AD paradigm in which subjects were allowed to give either or both associates as responses, Barnes and Underwood (1959) also predicted and found that longer retention intervals led to recovery from "unlearning" of the original association. Unlike previous studies, which allowed only one response, this new "modified modified free recall" procedure allowed Barnes and Underwood to assess absolute rather than just relative recovery of the original association.

Summary for Strategic Effects Account

To summarize, it is clear that strategic effects are often significant contributors to the misinformation effect. Whether a misinformation effect is detected and the proportion of the effect that is due to memory impairment depend on a number of factors, including the kind of test employed, overall memory for the original information, the delay between study and test, and the kind of materials used.

Although results from the modified test and tests of recognition are usually supportive of strategic effects as a partial account, the presence of a misinformation effect, even when these influences are controlled, suggests that memory impairment is also a fundamental component of the effects. In the comparison of the relative contributions of strategic effects and memory impairment, it is noteworthy that strategic effects do seem to account for a large part of the effects seen in the early experiments. In these experiments (see, e.g., Greene, Flynn, & Loftus, 1982; Loftus, 1979; Loftus et al., 1978), the subjects in the misled condition were generally 30–40% less accurate than were the control subjects (considering the relative proportions of correct responses). Later experiments designed to eliminate possible strategic effects and task demands sometimes have found differences in the 10–20% range (see, e.g., Belli, 1989; Belli, Lindsay, Gales, & McCarthy, 1994; Chandler, 1989), whereas other studies found very large effects, even when strategic factors were minimized (Roediger, Jacoby, & McDermott, 1996). This suggests that the strategic effects interpretation does not offer a complete explanation for the misinformation effect. But often it does explain a substantial portion of the effect, and it has been instrumental in improving the experimental control in this area of research.

THE BLOCKING ACCOUNT

An alternative explanation for the misinformation effect is the blocking hypothesis, which assumes that intervening exposure to the incorrect information impairs access to the correct information (see, e.g., Bekerian & Bowers, 1983; Bowers & Bekerian, 1984; Chandler, 1991; Morton, Hammersley, & Bekerian, 1985). The term *blocking* has sometimes been used loosely to refer to any multiple-trace hypothesis. Defined in that way, SAC's activation-based account of the misinformation effect would be one form of the blocking hypothesis. We prefer a narrower definition of the blocking hypothesis—as the subset of response competition hypotheses that involves the assumption that, at the time of the test, traces for both the original and the misleading items exist in memory and that the more recent trace blocks access to the earlier trace. According to this more restrictive definition, blocking is really an account of retroactive interference. This narrower definition is more in keeping with the common meaning of the term *blocking* and preserves important distinctions among different types of multiple-trace accounts that are sometimes lumped together.

The Headed Records Model

One of the most well-specified forms of the blocking hypothesis is based on the headed records model of memory (Morton et al., 1985). Briefly, the headed records model assumes that memory is organized as a system of records, each of which has a heading. Each heading is a brief description of the contents of its corresponding record. The critical assumptions are that people do not have direct access to the headings, and that records can only be accessed through headings. Records are accessed by matching a description (a combination of intentional search strategy and retrieval cues in the environment) to the headings. The model assumes a great deal of overlap in the content of records. Thus, multiple records may contain the same information but perhaps with different organization and different degrees of completeness. The framework assumes that, if multiple headings match a retrieval cue, the most recent one is given precedence for the response. Thus, retrieval errors occur when the description is not precise enough to be matched only to the proper heading but also matches another heading. In the case of the misinformation effect, because this other heading was created more recently and matches the retrieval cues, the wrong record is retrieved.

The headed records account has similarities to some aspects of the activation-based account. Where headed records assumes multiple records with greatly overlapping content, the activation-competition model uses the same nodes in the representation of multiple episodes. Thus, rather than storing each episode in a separate record, activation-based models represent multiple episodes by creating separate connections between nodes shared by multiple episodes and nodes representing unique aspects of individual episodes (see Figure 1). Although headed records and activation-based models are similar in some

respects, there are important differences, which will become clear in the discussion of issues on which the two models make divergent predictions.

Evidence Supporting the Headed Records Model

As evidence for the headed records model, Morton et al. (1985) point to Bekerian and Bowers's (1983) finding that the misinformation effect disappeared when the final memory test proceeded in the same order as did the sequence of events in the slides rather than in random order. Their argument is that, when the test follows the original order of events, information about sequence is included in the description used to access memory, and it is this sequential information that distinguishes the information presented in the slides from the information in the misleading questions.

The lack of an effect when the test questions follow the order of the original slide presentation would be inconsistent with the overwrite account, but it is not clear that the effect is robust; Bekerian and Bowers's (1983) null effect may be due to insufficient power rather than being a true null effect. Consistent with this interpretation, we know of no other experiments showing that the misinformation effect goes away when the test proceeds in chronological order, and at least one attempt to replicate the disappearance of the effect has failed (McCloskey & Zaragoza, 1985a). Without further corroboration, it is unclear whether the headed records model's prediction of this null effect is supported by the data.

Effects of Interpolated Tests: Further Support for Blocking

A class of studies using the modified test of McCloskey and Zaragoza (1985a) yielded results consistent with both the blocking account and the activation-based account. These studies used a modified recognition test, but subjects were given an interpolated test before it. In one of the studies, the interpolated test was a two-alternative forced-choice test with only incorrect responses as alternatives (Schooler, Foster, & Loftus, 1988); in the other, subjects made a yes/no decision about the misleading information (Belli, 1993).

In the Schooler et al. (1988) study, rather than an intervening narrative or misleading questions, an interpolated test was the vehicle that provided the misleading information. First, the subjects viewed a series of slides depicting a household burglary. They then performed an unrelated filler activity and read a narrative, just as in the standard paradigm. The difference was that the narrative contained no information about the critical items. Instead, after the narrative, the subjects were given an interpolated "recognition" test on which they were forced to choose an incorrect alternative (e.g., if the slides had shown *Ivory* soap, the alternatives on the interpolated test would be *Dawn* and *Palmolive*). Schooler et al. (Experiment 2) compared items for which only incorrect alternatives had been given on the interpolated test to items that had not appeared on the interpolated test at all. The results indicated

that committing to an incorrect alternative on the interpolated test hindered the subjects' recognition of the correct information on the final test, even though the alternatives from the interpolated test were not included on the final modified test.

These results are consistent with the blocking account's assumption that the more recent trace has precedence over older traces. To elaborate with the example introduced above, if *Ivory* was in the slides and *Dawn* and *Palmolive* were then presented on the interpolated test, these more recent traces would block access to the original information on the final test. Because subjects cannot access the original trace, they are more likely to have to guess on the final test (e.g., *Ivory* vs. *Sunlight*).

In an activation-based framework, the Schooler et al. (1988) results can be explained by the greater competition (fan) in the interpolated test condition (see Figure 3). The interpolated test condition adds additional traces to memory, whereas the no interpolated test condition does not. On the final modified test, some of the activation that would have otherwise gone to the correct response instead goes to these additional traces. Thus, the strength of the correct trace is reduced, relative to the strength of the corresponding trace in the no interpolated test condition.

In contrast to the Schooler et al. (1988) finding of impairment effects resulting from an interpolated test, Belli (1993) used a yes/no interpolated recognition test and concluded that the interpolated test yielded no impairment on a later modified recognition test. In Belli's (1993) study, subjects viewed a series of slides and read a postevent narrative containing either neutral or misleading information about some of the items in the slides—for example, a coffee can. In the slides, the brand of coffee may have been Folgers, in which case the narrative provided either neutral (*coffee can*) or misleading (*Maxwell House*) information about the critical item. The subjects were then given a yes/no interpolated recognition test on the information from the slides. The critical items on the interpolated recognition test were those that had been suggested in the misleading narrative (e.g., Maxwell House), so the correct answer was always "no," because the item had not appeared in the slides. Thus, unlike Schooler et al.'s, Belli's (1993) subjects were not forced to choose an incorrect alternative. Finally, the subjects were given a modified recognition test, on which they were asked to choose whether the original item (Folgers) or a completely new item (Nescafé) was shown in the slides. The item suggested in the narrative (Maxwell House) was not an option on the final test.

Belli (1993) was interested in the subjects' performance on the final modified recognition test, particularly on those items for which their responses had been consistent with the information from the narrative rather than the slide show. Table 1 summarizes the relevant results from Belli's Experiment 2. As expected, the subjects were less accurate on the interpolated test when they had been misinformed by the narrative than when they had received neutral information (62.5% vs. 32.5% acceptance of Maxwell House). Earlier studies had found similar results for

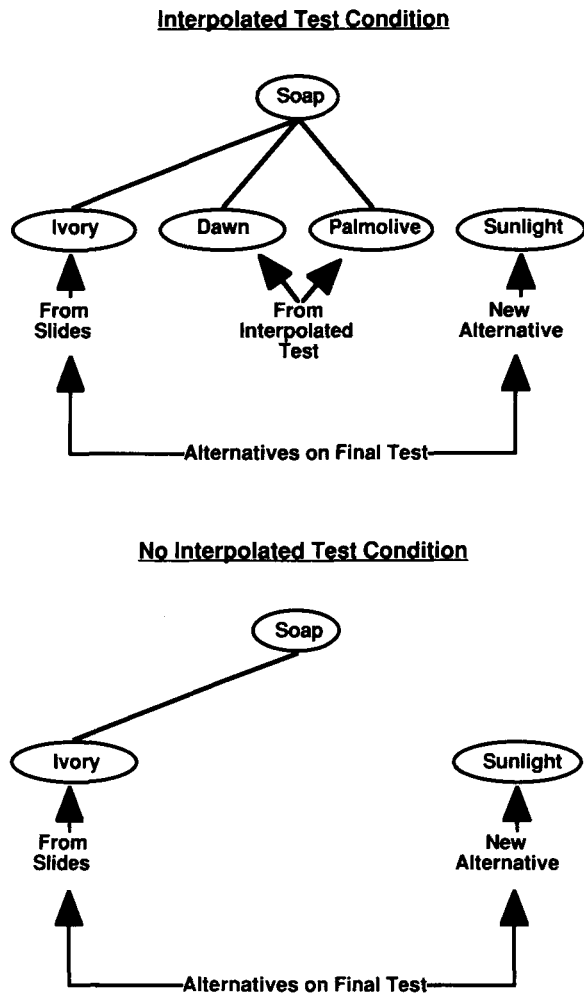


Figure 3. Alternatives on final modified test and memory structure for relevant concepts assumed by activation-based models for Schooler et al.'s Experiment 2. This figure is a simplified version of the representation, showing only the nodes and links critical in understanding the result of interest. In an activation-based model, this representation would be imbedded in a representation more like that in Figure 1.

yes/no recognition tests (Belli, 1989; Tversky & Tuchin, 1989). The critical issue with regard to the blocking hypothesis is the subjects' performance on the final test, conditionalized on their acceptance of the incorrect response on the interpolated test (emphasized cells).

As Table 1 indicates, the subjects in Belli's (1993) misled condition were no less accurate on the final test when they had accepted the misinformed item on the interpolated test (60.0%) than were subjects in the control condition who accepted this new alternative that had not been suggested previously (61.5%). Belli (1993) interpreted this null result as indicating that the act of accepting misinformation produced no impairment on the later recognition test. He argued that, if there had been a blocking effect resulting from misinformation acceptance, misled subjects who had accepted the misinformation in the recog-

nition test should have chosen the correct alternative on the final test less often, in comparison with subjects in the control condition who had accepted a completely new item. On the basis of this null effect, Belli (1993) argued that the blocking effect applies only to situations in which the incorrect information is available as an option on the test (e.g., Folgers vs. Maxwell House).

An alternative interpretation is that, if the memory representation for Folgers was so weak in the control condition that subjects could not reject Maxwell House (even though it was never mentioned), then the act of accepting Maxwell House on the interpolated test had the same interfering effect as an interpolated narrative. Instead, subjects added Maxwell House to their representation of the sequence of events. The final test then pits the relatively weak trace for the original information against a new alternative.

Note that the probability of having a relatively inaccessible link to Folgers is greater in the misled condition than in the control condition and is reflected in the acceptance rates for Maxwell House, the misleading item from the narrative (62.5% vs. 32.5%, respectively). However, even in the control condition, the subjects failed to reject Maxwell House, indicating that the Folgers link was weak, generally. This link was made even less accessible by the introduction of the misinformation. In short, it should be expected that both groups of subjects who accepted an incorrect alternative on the first test should perform equally on a later modified recognition test; the interpolated test essentially exposed *both* the control group and the misled group to misinformation. A more apt control would have been a condition in which subjects did not have an interpolated test. Belli's (1993) data suggest that accepting an incorrect alternative after reading misleading information does not impair access to original memory any more than does accepting an incorrect alternative following neutral information. They do not, however, speak to the issue of whether multiple tests produce greater memory impairment.

Roediger and McDermott (1995) provide further support for the argument that impairment effects from intervening tests are real and do not depend on the presence of the misinformation at the time of the final test. In that study, subjects were presented with word lists that included thematically related words. After seeing this list, there was a final recognition test, but prior to that test, one group was asked to recall the presented words, whereas another group performed an arithmetic task. Those who had free-recalled the list were much more likely to believe the prototype had been presented. This occurred because there were many intrusions of the prototypic word on the recall test and the subjects misattributed these intrusions as evidence of an actual presentation.

Roediger et al. (1996) also examined the effects of multiple recall tests on subjects' likelihood of reporting that they remembered seeing in the slides details that had appeared only in a misleading narrative. In their Experiment 1, the subjects first viewed a series of slides depicting a robbery. After seeing the slides, they read a mock police report, which was either accurate (control group)

Table 1
Percent and Number of Subjects Correct and Wrong on Interpolated Recognition and Final Modified Tests

| Presentation | | | Test | | | | | | | |
|--------------|---------|---------------|--|----|-----------------------------|----|--|----|--|----|
| | | | Interpolated Recognition (Yes or No to Maxwell House) | | | | Final Modified (Folgers vs. Nescafé) | | | |
| | | | % Accepted Maxwell House | | % Rejected Maxwell House | | % of Those Accepting Maxwell House Who Chose Folgers | | % of Those Rejecting Maxwell House Who Chose Folgers | |
| Condition | Slides | Narrative | N | N | N | N | N | N | | |
| Control | Folgers | Nothing | 32.5 | 39 | 67.5 | 81 | 61.5 | 24 | 69.1 | 56 |
| Misled | Folgers | Maxwell House | 62.5 | 75 | 37.5 | 45 | 60.0 | 45 | 75.6 | 34 |

Note—From “Failure of Interpolated Tests in Inducing Memory Impairment With Final Modified Tests: Evidence Unfavorable to the Blocking Hypothesis,” by R. F. Belli, 1993, *American Journal of Psychology*, 106, Table 2, p. 417. Copyright 1993 by Board of Trustees, University of Illinois. Adapted with permission.

or contained misinformation about some of the details in the slides (experimental group). The subjects were then given a cued recall test, with different instructions for different items on the test. For some of the items, the subjects were allowed to respond with an answer from either source (slides or narrative). For other items, the subjects were instructed to report only items that had appeared in the slides (slides only). Half of these slides-only items were tested under conditions of divided attention. Subjects were not forced to generate answers; they could elect not to respond to any or all items. Two days after the first session, the subjects returned for a second test, which was identical to the first day’s test, except that (1) they were instructed to respond to all items on the basis of their memory of the slides only, (2) there was no divided-attention manipulation, and (3) subjects also made remember/know judgments (Tulving, 1985) for each item they recalled.

Roediger et al. (1996) found a sizable misinformation effect on the first cued recall test. Recall of narrative items by misled subjects in the slides-only condition was 33%, as opposed to 10% in the control condition. Addressing the issue of impairment effects from intervening tests, Roediger et al. found that reporting the suggested information on the original cued recall test led to still larger misinformation effects on the second cued recall test. Most importantly, the conditional probability of recalling a suggested item on the second test, given that it was produced on the first test, was nearly identical, regardless of whether the item had originally been with either-source instructions (.73) or slide-only instructions (.76), and was considerably higher than the conditional probability in the control condition (.20). Roediger et al. argued that presumably, for some of the either-source items, the subjects knew on Test 1 that the item they reported came from the narrative. These results suggest that merely generating the suggested item led to later acceptance of it, *even though, at the time of the first test, the subjects had sometimes known that they were reporting on memories from the narrative.*

Roediger et al.’s (1996) results are consistent with a blocking interpretation in which recall of the narrative item on Test 1 blocks access to the original item on Test 2. These results are also consistent with the SAC perspective. Given that the item from the narrative was stronger on the

first test, it should also have been stronger at the time of the second test, particularly because recalling it on the first test gave it additional strengthening, whereas the item from the slides continued to decay.

Distinguishing Blocking and Activation-Based Accounts: The Reversed Design

Although the previous experiments do not distinguish among possible multiple-trace interpretations of the misinformation effect, the *reversed design* does differentiate blocking from other multiple-trace accounts. In this design, subjects first are given the misinformation, then see the slides, and finally are given a yes/no recognition test on the slides. This procedure yields memory impairment effects nearly as large as those for the standard order of events (Lindsay & Johnson, 1989b; Rantzen & Markham, 1992).

Precedence of the most recent trace is a critical assumption of the blocking account. If subjects base their responses solely on recency, one would predict that the misinformation effect should go away for the reversed design, because the most recent information is the information from the slides (the correct information). Results from reversed design experiments also raise problems for the trace alteration account. If the misinformation effect occurs because a new trace overwrites a preexisting trace, the reversed design should eliminate the misinformation effect because, at the time the misinformation is encoded, there is nothing for it to overwrite. The data are inconsistent with this prediction of the trace alteration account.

Whereas the trace alteration and blocking hypotheses have difficulty accounting for this result, the finding is consistent with an activation-based account, such as SAC. According to SAC, the trace representing the misinformation seen earlier in the experiment is still somewhat available, as is the correct information. If the misinformation is presented as a probe and subjects base their judgments primarily on the activation of the trace for the test item without retrieving the associated source information, they might false alarm on the critical items. Evidence from other paradigms supports the notion that subjects sometimes make judgments on the basis of cue familiarity, rather than on the basis of careful search (Metcalfe, Schwartz, & Joaquim, 1993; Reder & Ritter, 1992).

Summary for Blocking Account

With the simple assumption that the most recent trace takes precedence and blocks access to older traces, the blocking account adequately explains impairment effects in the original studies (see, e.g., Loftus et al., 1978), in studies that use a modified recognition test (see, e.g., Chandler, 1989), and in studies that employ interpolated tests (Schooler et al., 1988). Although adequate for explaining those classes of experiments, blocking cannot account for results from studies using the reversed design, in which the most recent trace is actually the correct information. Any framework that hopes to explain the misinformation effect must be sufficiently well specified to account for impairment effects, regardless of when the misinformation is encountered. The headed records model or other instantiations of the blocking account may be able to overcome this shortcoming by positing different assumptions, but to our knowledge none has done so to date.

With regard to memory impairment effects, we see the original design as a variant of the traditional retroactive interference paradigm (see, e.g., McGeoch & McDonald, 1931; Osgood, 1949) and the reversed design as a variant of traditional proactive interference studies (see, e.g., Greenberg & Underwood, 1950; Keppel & Underwood, 1962; Underwood, 1957). Any model of memory that accounts for these two effects probably also accounts for the majority of the memory impairment effects discussed so far.

THE SOURCE-MONITORING ACCOUNT

In an attempt to give a more complete explanation for the existing and emerging results, Lindsay and Johnson (1987) suggested that whether a misinformation effect occurs depends on variables that affect subjects' likelihood of retrieving the correct source information before responding. There is extensive evidence from a variety of paradigms that suggests that people sometimes have difficulty determining the source of their memories (for a review, see Johnson, Hashtroudi, & Lindsay, 1993).

According to Lindsay and Johnson (1989a), one way source misattributions can occur is for the source of the original information not to exist in memory, either because it has decayed away or because it was never encoded in the first place. In these cases, people try to infer the source of their memory, and they generate a plausible but incorrect source. A second way source misattribution errors could arise is for the subjects to tend to respond on the basis of familiarity with the cue instead of with the source information. In many paradigms that are used to investigate the misinformation effect, either the misinformation is more familiar than is the original information, or the misinformation is the only alternative presented. Thus, it is accepted because of its relative familiarity at test. Johnson et al. (1993) enumerated several other variables that influence the likelihood that people will make source-monitoring errors. These mediating factors include the extent to which subjects imagine visual or auditory features as they encode verbal misinformation, the extent to which

the misleading information is central rather than tangential to the series of events, and the extent to which subjects are stressed or distracted during the encoding of the misinformation or during the test.

It is important to emphasize that *source misattribution* and *misinformation effect* are not interchangeable terms. Source misattributions are situations in which subjects erroneously attribute the source of a remembered or recognized item to one source when it actually came from another. In principle, source misattributions are possible both in control conditions and in misinformed conditions. A misinformation effect occurs when there is a difference between misled and control conditions on recall or recognition accuracy. Thus, misinformation effects may be the result of source misattributions that are made in the experimental condition but not in the control condition, but the two effects are theoretically independent.

As an illustration of how source judgments might be made on the basis of familiarity rather than on the basis of retrieval attempts, consider Lindsay and Johnson's (1989a) explanation for the results from Loftus et al.'s (1978) original testing procedure. Because most of the questions on the test were about items seen in the original event or distracter items that were never presented in any form, they could be answered correctly simply by judging which of the choices at test seemed most familiar. For control items, this strategy yields accurate performance; the items presented in the slides are generally more familiar than are completely new items. On the critical items, however, the item from the narrative has been encountered more recently, so it is more familiar than is the one from the slides and is chosen on the test.

Yes/no versus source-monitoring tests. Lindsay and Johnson (1989a) gave subjects one of two types of tests. In the yes/no recognition test, the subjects studied a slide of a cluttered office and then read either a completely accurate description of the slide or a description containing misinformation. They were then given a list of items and were instructed to check "yes" or "no," depending on whether or not the item appeared in the slide. In the source-monitoring task, instead of checking "yes" or "no," the subjects were instructed to check whether the item previously appeared in the text only, in the picture only, in the text and in the picture, or not at all. Compared with the subjects in the yes/no condition, the subjects in the source-monitoring task were much more accurate at discriminating items that were in the slides from those that were not. Zaragoza and Lane (1994) also compared yes/no and source-monitoring tests, using a typical slide-encoding procedure, and obtained similar results.

According to the source-monitoring account, if subjects are in a task that encourages them to inspect the source information before responding, there should be a reduction in the misinformation effect to the extent that the source information can be retrieved. The source-monitoring explanation for Lindsay and Johnson's (1989a) results is that the yes/no task is amenable to rapid, familiarity-based judgments; in contrast, subjects in the source-monitoring

condition are more likely to attend to source information because they are explicitly instructed to do so.

The activation-based account is largely consistent with this explanation. In SAC, part of the representation involves an association of encoding context (source) to the encoded facts. Whether the contextual information is accessed depends on (1) the strength and availability of the context (source) and (2) the effort (time) subjects are willing to allocate to retrieval of that link. As in the source-monitoring account, we assume that subjects in the yes/no task will be less accurate than subjects specifically asked to judge the source information. A misinformation effect is found in the yes/no task because, at test, activation spreads from the queried item and the schema or scenario, intersecting at an encoding event node, which passes over threshold (see Figure 1). If they get intersection of activation, they say "yes," if not, they say "no," without ever inspecting the source information associated with the event node.

Lindsay and Johnson (1989a) argued that yes/no judgments are probably made more quickly than are source-monitoring judgments. SAC makes the same prediction; yes/no judgments should be made more quickly than source-monitoring judgments because it takes additional time to retrieve and inspect associated source information. The instructions for the source-monitoring task require that the source information be active before a response is given (unless the subject is guessing because no information can be retrieved). Thus, to the extent that the source information can be retrieved, subjects in the source-monitoring condition should be less likely to falsely classify an item as being from the slide. One prediction that stems from this is that, if a delay is built into the yes/no judgment such that subjects could not respond until as long after the presentation of the test item as it takes for source-monitoring subjects to respond, yes/no subjects will be less likely to say that a suggested item was in the slide. We make this prediction because the added delay allows for more time to inspect the nodes representing contextual information.

According to the activation-based account, the difference between the yes/no and source-monitoring tasks is due to a difference in response mode adopted by the subjects. We should clarify that the model has no representation of *response mode* and does not itself predict which mode subjects will adopt. Although RTs were not recorded, it is likely that the subjects in the source-monitoring task took longer to respond to the questions than did the subjects in the yes/no task. Again, this prediction is the same as that made by Lindsay and Johnson (1989a). By changing the task from yes/no to source monitoring, one is forcing subjects to conduct a more effortful match involving more elements, some of which send little activation to the schema by virtue of their high fan. This difference is reminiscent of the Moses illusion (Kamas, Reder, & Ayers, 1996; Reder & Kusbit, 1991), in which subjects responded to questions much more quickly if they only had to access gist information rather than to carefully inspect the question for possible tricky elements. Thus, rather than changing the process, the source-monitoring task can be thought of as just asking for more information from the represen-

tation. We suspect that, if subjects were required to classify the source of the information before a deadline, their performance would become more similar to that in the yes/no condition. Further evidence for this interpretation is provided by Kelley and Jacoby (1996), who employed a response deadline in a memory test for recognition of words from one of two previously presented lists. When the deadline was very short, subjects were more likely to erroneously accept words from the to-be-excluded list, presumably because the enforced fast responding precluded recollection of the source list information.

Thus far there are no obvious differences between the source-monitoring and activation-based accounts, except that the activation-based account incorporates the assumptions of source monitoring in a memory model. The source-monitoring account is not offered as a structural model of the memory. Instead, it is offered as a description of people's behavior in some cases, depending on the demands of the particular task. We now turn to an experiment that differentiates between these two accounts.

Manipulating Discriminability of the Original and Misleading Information

A second type of evidence used to support the source-monitoring account comes from an experiment that manipulated discriminability of the sources of correct and misleading information (Lindsay, 1990). This experiment had two conditions: high discriminability and low discriminability. In both conditions, the subjects saw a slide sequence and listened to an accompanying recorded description of the events. After the slide and audio presentation, both groups listened to a second recorded description of the events. In the low-discriminability condition, this second recorded description immediately followed the slide show, and the recording was in the same voice that accompanied the slides. High-discriminability subjects heard the second verbal description 48 h later in a different part of the room and in a different voice than the one they experienced with the slides. Low-discriminability subjects were instructed to imagine each object as it was described in the second description; high-discriminability subjects were instructed to mentally repeat each word they heard. Forty-eight h after the slide show (i.e., immediately after the second recording for the high-discriminability group), both groups were given a cued recall test for the three control items and for the three items about which misinformation had been given in the postevent information. In accord with a methodology that was similar to the Jacoby, Woloshyn, and Kelley (1989) *logic of opposition* paradigm, subjects in both conditions were explicitly instructed that any information they remembered from the narrative—that is, information from the recording that was played separately from the slide presentation—was wrong. Thus, they were implicitly instructed to attend to the source of their responses.

The data from Lindsay's (1990) experiment are summarized in Table 2. As can be seen in the table in the column labeled *Suggested*, a source misattribution effect was obtained only in the low-discriminability condition, and

Table 2
Mean Proportion of Responses of Each Kind by
Acquisition Condition and Item Type

| Condition | Recall Response | | |
|---|-----------------|---------|-------|
| | Suggested | Correct | Other |
| High discriminability | | | |
| Control items | .10 | .48 | .36 |
| Misled items | .13 | .39 | .39 |
| Low discriminability | | | |
| Control Items | .09 | .51 | .27 |
| Misled items | .27 | .45 | .23 |
| Difference Scores (High Minus Low Discriminability) | | | |
| Control | .01 | -.03 | .09 |
| Misled | -.14 | -.06 | .16 |

Note—From “Misleading Suggestions Can Impair Eyewitnesses’ Ability to Remember Event Details,” by D. S. Lindsay, 1990, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, Table 1, p. 1081. Copyright 1990 by American Psychological Association. Adapted with permission.

not in the high-discriminability condition (3% effect vs. 18% effect). This result is reminiscent of Postman and Underwood’s (1973) finding that subjects’ ability to discriminate list membership diminishes over time. Besides replicating this result, Lindsay’s data also suggest a theoretically important dissociation: although only the low-discriminability condition led to more source misattributions, overall accuracy (proportion correct) and the degree of memory impairment did not differ for low- as opposed to high-discriminability groups.

Lindsay (1990) explained this dissociation between the source misattribution effect and memory impairment in terms of source discriminability—that is, in the low-discriminability condition, the subjects were less able to determine the source of their memory, so they reported the interpolated item more often. In the high-discriminability condition, the subjects recognized that the source of the interpolated item was the narrative, so they did not report it. Despite the differences in source discriminability, the strength of the original information is the same in both conditions, so no difference is found in terms of memory impairment. This dissociation is consistent with Lindsay’s (1990) claim that source discriminability and trace strength are independent constructs.

The activation-based account is similar to the source misattribution explanation, instantiating the differences in availability of source information in terms of the activation levels of the links to the source information (e.g., the links from *stop* and *yield* to the *context* nodes in Figure 1). In the high-discriminability condition, the misinformation has a stronger memory trace (because it is more recent) but so does the episodic information associated with it. Because subjects inspect these context nodes before responding, they do not give the misinformation as a response, and no source misattribution effect is found in the high-discriminability condition. In contrast, in the low-discriminability condition, both the original and the misleading information’s traces have weakened over the 2-day delay between misinformation and test. Over this delay, the context nodes

that allow one to infer source have also decayed. At the time of the test, the residual strength of the two traces are roughly equivalent, so the subject has nothing to go on except the fact that both the original and misleading concepts are related to the overall schema.

Besides explaining this basic result, the activation-based account also makes correct predictions for the remainder of the data Lindsay (1990) reported. Columns 2 and 3 of Table 2 indicate the proportion of correct responses and of responses that were neither the suggested item nor the correct response.

To address Column 2 first, both low- and high-discriminability subjects were less often able to recall the items from the slides for misled items than for recall control items. Thus, even though there was no effect in the high-discriminability condition in terms of likelihood of reporting the misinformation, the misinformation *did* have an effect. Lindsay (1990) argued that this is evidence for true memory impairment *even when subjects accurately recall the misinformation’s source*. In terms of establishing the existence of true memory impairment (as opposed to response biases or strategic effects), this is an important point.

Consistent with Lindsay’s (1990) conclusion, activation-based models such as SAC predict impaired recall of the original information for misled items in comparison with control items, regardless of whether the misinformation’s source is also retrieved. In both low- and high-discriminability conditions, the introduction of the misleading information in the narrative adds an extra association to the critical node in memory (e.g., the *sign* node in Figure 1). When asked to recall the type of sign, activation must be shared by both associations, meaning that less activation reaches the correct item in the misled condition than in the control condition, for which only the original association exists. The reduced activation to the original item translates into a reduced likelihood of generating it on the final test.

Column 3 of Table 2 shows the mean proportions of *other* responses for the control and misled conditions for low- and high-discriminability conditions. These are trials on which the given response had appeared neither in the slides nor on the audio recording. Overall, subjects in the high-discriminability condition were more likely to generate a new item than were their peers in the low-discriminability condition. Although the interaction was not reliable, this effect was especially pronounced for items about which the subjects had received misinformation (16% for misled items vs. 9% for control items). The activation-based interpretation would explain this difference as resulting from interference of the misleading information. As discussed above, in the low-discriminability condition, the information from the narrative is given as a response on the test because of source confusion, yielding a source misattribution effect (Column 1); however, in the high-discriminability condition, the misinformation’s source is salient, so they inhibited reporting it. Instead, the interference from the narrative has its effect by making the correct information

less available. The additional association's presence in memory still interferes with the subjects' ability to retrieve the correct answer, which leads to an increase in *other* responses, as compared with the low-discriminability condition. Simply put, the effect of the misinformation in the low-discriminability condition is to produce source misattributions for the suggested items; the effect of the misinformation in the high-discriminability condition is to produce more intrusions.

Lindsay (1990) explained the final column of Table 2 (responses that were neither correct nor the suggested item) in terms of a shift in the subjects' willingness to guess. He explained the lower guessing rate in the low-discriminability condition in terms of a metacognitive strategy whereby subjects were "cautious about reporting details that came to mind lest they err by reporting details from the postevent narrative" (p. 1082). His logic was that low-discriminability subjects were biased against guessing because the instructions stated that anything from the narrative was wrong. Were this the case, however, one would expect subjects to be more likely in the low-discriminability condition to refuse to answer at all. This is clearly not the case: low-discriminability subjects gave no answer 9% of the time, compared with 8% for the high-discriminability subjects. Thus, although this differential-guessing bias is one possible explanation for the difference in *other* responses, it is not strongly supported by the data. Moreover, as discussed above, the structural account involving activation explains this pattern without requiring additional assumptions about subjects' metacognitive strategies.

Strength of belief in false memories. An issue of recent inquiry in understanding the misinformation effect is subjects' qualitative experience when the effect occurs. Are some false memories phenomenologically the same as real memories? This question is important because it goes to the heart of the debate over whether the misinformation effect is a true memory phenomenon. Is the misinformation effect due to real belief that the suggested items were seen, or are subjects only really sure that they *experienced* the critical item somewhere in the experiment and, if pressed, would *not* report it as having been *seen*?

It may be that the manipulations commonly employed in these experiments led subjects to truly believe that they saw the suggested details. It could also be that these manipulations only made subjects more likely to *report* that they saw the suggested item, even though their memories for real and suggested items were phenomenologically different. Several forms of source-monitoring tests have been used in an attempt to determine whether the misinformation effect can be explained by subjects' failure to access source information at all or whether people genuinely confuse source information (Belli et al. 1994; Lindsay, 1990; Lindsay & Johnson, 1989a, 1989b; Roediger et al., 1996; Weingardt, Loftus, & Lindsay, 1995; Zaragoza & Koshmider, 1989; Zaragoza & Lane, 1994; Zaragoza & Mitchell, 1996). There is evidence that sometimes people are good at attributing source information, sometimes they are poor, and sometimes they do not access source information at all.

Findings by Zaragoza and Koshmider (1989) suggested that subjects' memory impairment in other studies may not reflect source misattributions but rather failure to access source information at all. This study consisted of a slide show depicting an office robbery, followed by a misleading narrative and then by a recognition test (also slides). The subjects were given four alternatives on the final test: saw, read, consistent, and inconsistent. The instructions were worded such that *saw* was the correct response, even if the subjects remembered both seeing and reading the item. *Read* was the correct response only if they remembered reading the item in the narrative *and* did not remember seeing it in the slides. *Consistent* and *inconsistent* referred to whether the item in the test slide corresponded to the subjects' memory for the robbery episode. These last two alternatives were to be used only for those trials on which subjects did not remember the source information at all. Zaragoza and Koshmider found that, both at no delay and at a 1-day delay between the misleading narrative and the recognition test, subjects were no more likely to report having *seen* the misleading items from the narrative than they were the control items they had not seen at all. However, the subjects were less likely to judge items from the narrative as being *inconsistent* with their recollections of the robbery episode than they were control items. Most previous experiments did not differentiate between responses for which subjects were confident of the source of their memories and responses for which subjects were less confident. Assuming that *consistent* judgments would have appeared as false recognitions or false recalls under more typical testing procedures, Zaragoza and Koshmider's results suggest that the observed misinformation effects in other studies may sometimes be due to acceptance of misinformation when the source cannot be retrieved rather than to misattributing the source of the misleading items as being the original slides. This view is also supported by Loftus and Hoffman (1989).

Although Zaragoza and Koshmider's (1989) data support the notion of misinformation acceptance when a source is ir retrievable rather than the notion of a true belief that the suggested item was actually seen, the evidence is not conclusive. The assumption that *consistent* judgments would end up as false recollections or recognitions under testing procedures is purely speculative without a direct comparison of these conditions within a single experiment. It could be that the *consistent* option drew upon responses that would otherwise have been low-confidence rejections of the suggested information.

Effects of encoding task. Recent studies by Zaragoza and her colleagues strongly suggest that some subjects truly believe that they *saw* items that were only suggested to them in a written narrative (Mitchell & Zaragoza, 1996; Zaragoza & Lane, 1994; Zaragoza & Mitchell, 1996). These experiments that were designed to determine subjects' phenomenological experience differ from previous attempts (e.g., the logic of opposition) in two important ways. First, subjects' responses were not constrained by a set of logical rules that could contradict the subjects' phenomenological experience. The subjects were allowed to

say whether they thought a test item came from the original slide presentation, from the narrative, or from both the slides and the narrative or whether they thought it was new to the experiment. The second difference from most earlier experiments was that the response procedure allowed the subjects to express uncertainty rather than forcing them to choose an alternative. The subjects were allowed to express—on a 7-point scale that ranged from *Definitely Yes* to *Definitely No*, with *Unsure* in the middle—whether each test item appeared in the slides and whether it appeared in the narrative.

Zaragoza and Lane (1994) used a procedure in which the critical information in the narrative *supplemented* rather than contradicted the information from the slides. For example, instead of the narrative suggesting a different instance of an item from the same category as an item from the slides (e.g., a stop sign when the slides had shown a yield sign), the suggested item in this paradigm was plausible but never appeared in the slide at all (e.g., a coat rack in the room when the slides had not contained one). The advantage of this design is that it eliminates situations in which subjects notice a contradiction between the slides and the misleading interpolated task and then possibly make faulty inferences about the veridicality of their original memories. By reducing the likelihood of these extra inferences, Zaragoza and Lane can more confidently conclude that any effects are the result of memory effects rather than the result of inductive processes.

In a series of experiments, Zaragoza and Lane (1994) manipulated the extent to which subjects were likely to actively reconstruct the original series of events during the intervening task. This manipulation took the form of varying the intervening task. Some subjects read an intervening narrative that suggested items that were not in the slides, as is the common procedure in many misinformation effect experiments. Others answered questions that presupposed the presence of objects that had not appeared in the slides. Still other subjects received the sentences from the narrative in scrambled order and were asked to reconstruct the original order of events from their memory of the slides. The critical difference in these three intervening activities was that answering questions and unscrambling the narrative *required* retrieval of the information contained in the slides, whereas reading the narrative did not.

The results indicated that, when the intervening task required an active mental reconstruction of the original series of events (the questions and the scrambled narrative), subjects were more likely to make source misattribution errors than they were when they simply read the narrative. In an examination of the qualitative aspects of the source misattributions people made, subjects were often confident in their misattributions. Zaragoza and Lane's (1994) Experiment 3 used a 7-point Likert-type scale, ranging from *Definitely yes* to *Definitely no*. The subjects answered *Definitely Yes*, *Saw in slides* to 57% of their misattributions, compared with only 35% of their misattributions for control items, supporting the view that subjects were highly confident in their source misattributions for the suggested items.

Interestingly, for most source misattribution errors, the subjects attributed the narrative-only items to *both* the slides and the narrative. Less than 11% of the time did the subjects claim to have seen but not read about an item that had only been suggested to them. As a proportion of source misattributions, for over two thirds of the source misattributions, the subjects said they both saw and read the item; for less than a third of the source misattributions did they claim just to have seen an item that actually they only had read in the narrative.

On the basis of the subjects' relatively high confidence in their misattributions, it appears that source misattributions are not simply guesses in situations where source information cannot be retrieved. More often than not, when errors are made, subjects truly believe they *saw* the suggested items. Similar results have been obtained by Zaragoza and Mitchell under different conditions, indicating the robustness of the phenomenon (Mitchell & Zaragoza, 1996; Zaragoza & Mitchell, 1996).

Zaragoza and Lane (1994) concluded that whether subjects believe they *saw* the suggested item is determined, at least in part, by the extent to which they mentally reconstruct the original episode during the intervening task. They suggested as one possible explanation that, in reconstructing the original episode, subjects are more likely to use imagery than they are when the intervening task is more passive. Having created a mental image during the intervening task, subjects might then sometimes misattribute the source of the image they retrieved during the final test to the slides, when the trace for the retrieved image had actually been created during that intervening task. There is abundant evidence to suggest that people are sometimes unable to accurately distinguish between memories for real and imagined events (see, e.g., Garry, Manning, Loftus, & Sherman, 1996; Johnson, 1988; Johnson & Raye, 1981). Although no direct manipulations of imagery instructions were used, this explanation seems reasonable, given the tasks Zaragoza and Lane used. The kind of memory misattributed—mental image or proposition—is irrelevant for most models of memory because the mechanism is the same in all cases.

The activation-based account is consistent with Zaragoza and Lane's (1994) interpretation of the cause of the effect: Source misattribution errors should largely come from situations in which processing the misleading information requires co-activation (or retrieval) of the originally presented information. When subjects reconstruct the original events during the interpolated task, they lay down memory traces for that reconstruction. These memory traces are highly similar to the traces laid down during the slide show, containing both propositional information about the series of events and relations and newly created representations of images. While piecing together the scrambled narrative, subjects must reconstruct the order of the events originally depicted. This reconstruction requires the reactivation of nodes that were created during the slide show, many of which represent visual images. During this reconstruction, subjects also sometimes create

a mental image of the suggested item, and they associate that image with the images created during the slide show. The representation of the suggested item then is of the same type as is that of the items that were actually seen and has many of the same contextual associations. There are some differences between the nodes in the original representation and this newly created node (e.g., different temporal associations), but compared to unimagined objects, the differences are small. At test, the cue reactivates the visual image created during study, and subjects infer that it was created during the original slide presentation, as were most of the other items for which they retrieve visual images. On the basis of this inference, they report with high confidence that it was on the slides. The cue also reactivates the contextual associations from reading and unscrambling the narrative, so subjects report that the item was in the scrambled narrative.

Repeated exposure to misinformation. The results from several studies suggest that repeated exposure to misinformation may increase people's likelihood of accepting false information (see, e.g., Hyman, Husband, & Billings, 1995; Loftus & Ketcham, 1994). It should be noted that, in the interest of ecological validity, these studies confounded repeated exposures with the possible demand characteristics and length of the delay between the original information and the final memory test. However, recent research has examined effects of multiple exposure to misinformation independent of these other variables. Repeated exposure to misinformation has been manipulated both at study (Mitchell & Zaragoza, 1996; Zaragoza & Mitchell, 1996) and at test (Roediger et al., 1996). In both cases, there was greater memory impairment following multiple exposures than following a single exposure to the misinformation.

Zaragoza and Mitchell (Mitchell & Zaragoza, 1996; Zaragoza & Mitchell, 1996) manipulated the number of times (one vs. three) and the number of modalities (one vs. three) in which subjects encountered the misinformation before being given a source discrimination test. They found that (1) with repeated exposure to suggestions, subjects were more likely to believe they actually saw items that were only suggested to them (but see Warren & Lane, 1995), (2) repeated exposure caused the effect to persist for at least 1 week, (3) presenting the suggestions in multiple modalities led to more source misattributions for both control items and suggested items, and (4) varying the spacing of multiple presentations had no effect on subjects' likelihood of making source misattributions.

As described earlier in the blocking section of this article, Roediger et al. (1996) examined the effects of repeated exposure to misinformation through multiple testing. Their subjects were given two cued recall tests, and the second test asked for items only from the slides. For each item recalled on the second test, the subjects made remember/know judgments (Tulving, 1985). For items that were tested twice, when the subjects reported seeing in the slides items that were only suggested by the narrative, they were more likely to *remember* than to *know* they had seen these items—that is, to misattribute source. Conversely,

for items that were not presented on Test 1, items spuriously recalled on Test 2 more often received *know* judgments than they did *remember* judgments. Furthermore, the proportion of *remember* judgments for suggested and twice-tested items was in the same range as that of recalled items that had actually appeared in the slides. These results suggest that repeated testing of misinformation made the subjects' phenomenological experience of accurate and false memories very similar. Assuming that recall requires a reconstructive process, the explanation for Roediger et al.'s remember/know data is essentially the same as above. During the first recall test, the subjects created a mental image of the item and then incorrectly inferred, on the second recall test, that the source of this image was the slides. For control items, there was no trace for a mental image, so there was no incorrect inference as to their source.

Overall, the evidence from these studies clearly suggests that repeated exposure to misinformation, whether at study or at test, increases the size of the misinformation effect and leads to more source misattributions, in comparison with a single exposure to misinformation (see also McDermott, 1996; Schooler et al., 1988). Furthermore, generating the misinformation can lead to source misattributions even when the correct source was known at the time the misleading item was generated. Phenomenologically, false memories for suggested items are often very similar to memories for items that were actually seen (see also Mitchell & Zaragoza, 1996; Payne, Elie, Blackwell, & Neuschatz, 1996; Read, 1996; Roediger & McDermott, 1995).

In terms of explaining why subjects make source misattributions under some conditions and not others, it appears that integration of the original memory with the interpolated information is an important factor. Experiments examining source misattribution that have used integrative interpolated tasks have found genuine source confusions (Roediger et al., 1996; Zaragoza & Lane, 1994). Experiments using procedures that do not encourage reintegration have not found evidence for source confusions; instead, there is the suggestion that, under these conditions, source misattributions may result from subjects' failure to access source information at all (Zaragoza & Koshmider, 1989). It remains to be seen whether the other effects discussed in this section, including greater misinformation effects with multiple modalities and effects of multiple narrative presentations (Mitchell & Zaragoza, 1996), are also the results of a greater propensity for subjects to integrate the suggested item with the representation of the original series of events.

Summary of Source-Monitoring Experiments

There are many similarities between the source-monitoring account and the activation-based account. It is fair to say that, in some cases, the activation-based account merely instantiates the assumptions of the source-monitoring account in a mechanistic model of memory. Source monitoring and SAC make the same predictions for the same reasons with regard to performance in yes/no

as opposed to source-monitoring test instructions and with regard to multiple presentations of misinformation. There are, however, a few important differences between the source-monitoring explanation for the misinformation effect and the activation-based account, as noted earlier. First, with regard to Lindsay's (1990) data on source discriminability, SAC correctly predicts aspects of the data not directly addressed by the source-monitoring account. Second, SAC also predicts a general effect in RTs, where the greater competition for activation in the misled condition should lead to increased RTs because activation must be divided among more links. Source-monitoring makes no predictions with regard to RTs.

In sum, we see nothing *wrong* with the source-monitoring account; to our knowledge, it makes no faulty predictions. Furthermore, some assumptions of the source-monitoring account are not instantiated in the SAC model, but are instead taken by SAC as given—for example, whether subjects are biased to make familiarity-based source misattributions rather than to make source-monitoring judgments. SAC is a model of declarative memory that can influence strategy choice (see, e.g., Reder & Schunn, 1996; Schunn et al., 1997). In some ways, the theory is impoverished in comparison with architectures, such as ACT-R, that instantiate procedural as well as declarative memory. It is a goal of these more complex architectures to handle biases in strategy choice (see Lovett & Anderson, 1996, for an example).

GENERAL DISCUSSION

The primary goal of this review has been to compare several theoretical accounts that have been proposed to explain people's memorial susceptibility to false information or misleading suggestions. Each of the explanations reviewed can account for some of the findings, with varying degrees of success. Each explanation is able to account for the particular set of data for which it was crafted, but the field lacks a plausible explanation for the corpus. Most theories make either no predictions or invalid predictions with regard to different methodologies within this arena. Trace alteration adequately accounts for the original data but is inconsistent with subsequent findings from the modified test and from studies using the reversed design. Blocking also adequately accounts for the original phenomenon but falsely predicts a null effect for the reversed design. The strategic effects account correctly predicts the null effects often found with the modified-testing procedure but cannot account for the presence of the effect under some conditions. The source-monitoring explanation adequately and accurately accounts for the presence or absence of an effect, depending on the methodology, but, without additional assumptions concerning the representation and processes of memory, it cannot account for other aspects of the data. In our view, an activation-based framework best accounts for the data across all the methodologies reviewed. We have tried to focus on some of the major issues and apparent inconsistencies in the lit-

erature, accounting for them using a relatively simple framework. The greatest appeal of this framework is that it accounts for many of the various misinformation effect phenomena with the same assumptions used to explain many other cognitive phenomena—for example, the fan effect, the feeling of knowing, the Moses illusion.

One instantiation of an activation-based framework that could account for the data reviewed here is the source of activation confusion (SAC) model of memory. SAC has much in common with other semantic network memory models, but it emphasizes a relatively unique assumption that makes it particularly well suited for the explanation of illusions of memory. SAC assumes that many such illusions or errors arise from the misinterpretation of activation.⁵ Specifically, people can respond to high activation on a node but often cannot know *why* it is high. According to the SAC model, many of the cognitive illusions and memory errors discussed in the literature can be accounted for by people's failure to identify the source of activation before giving a response or by their failure to achieve an exact match between an environmental source of activation and the correct memory representation.

SAC has not been used to formally model misinformation effect data, but it has been used to model other empirical data that conceptually seem to share many of the same memory mechanisms. For example, SAC has been successfully applied to the modeling of the mirror effect with word frequency, of continuous remember/know judgments, and of feeling of knowing—all involving issues of strengthening, decay, and interference. Undoubtedly, activation-based frameworks besides SAC (e.g., ACT-R) can also account for many of these data and disentangle apparently discrepant results using only a few, well-specified assumptions. One appeal of well-specified frameworks such as these is that, with the same sets of assumptions, they also account for many other cognitive phenomena (see, e.g., Kamas & Reder, 1994; Nhoyvanisvong & Reder, in press; Reder & Gordon, 1997; Reder et al., 1998; Reder et al., 1997; Reder & Schunn, 1996; Schunn et al., 1997). The general advantage of adopting an activation-based perspective in addressing the misinformation effect literature is that it makes clear predictions for more aspects of the data than does any other perspective offered to date, and it bases these predictions on the same set of assumptions that is used to account for other phenomena.

Activation-based models are well known for their ability to predict and account for the *fan effect* (see, e.g., Anderson, 1974b), which takes its name from the "fan" of associations from concept nodes, as shown in Figure 1. The fan effect refers to the reaction time interference effect analogous to interference effects observed in accuracy of recall or recognition in nontimed studies. Regardless of the dependent measure, there is the general finding that, the more associations one has made to a concept, the worse the performance (slower or less accurate) in accessing a single association.⁶ Activation-based models predict this result because, as the number of associates increases, the activation is shared by more associates. In SAC, less

activation translates into a lower probability of a concept going over threshold; in ACT-R, the mechanistic account is similar.⁷ The notions of less activation for each competing associate, strengthening because of repeated exposures, and decay of strength because of the passage of time have been used to model word frequency effects in recognition and remember/know data (Reder et al., 1997) and feeling-of-knowing (Reder & Schunn, 1996; Schunn et al., 1997). Many of SAC's explanations for the misinformation effect phenomena are based on the same sort of representation and assumptions; specifically, the presence or absence of extra associations to a concept can have important implications for memory, *even when those associations are not queried*.

To our knowledge, only one other formal model, CHARM, has been proposed to account for the misinformation effect (Metcalf, 1990). CHARM is a distributed memory model in which the interpolated item is *blended* or superimposed with the original trace. Thus, CHARM is a single-trace model, and it accounts for the misinformation effect by means of trace alteration. At the time it was proposed, there was some debate over CHARM's ability to account for the data (see Lindsay, 1991; Metcalf, 1991; Metcalf & Bjork, 1991; Schooler & Tanaka, 1991). Today the model seems less compelling than it did originally; it is not clear how the model would account for data showing that the impairment effect diminishes with time (see Windschitl, 1996, for a discussion) or for the small but reliable misinformation effect using the modified recognition procedure (Payne et al., 1994). Indeed, in all six of Metcalf's simulations of the modified test, using 1,000 observations per point, the misled condition was *more* accurate than was the control (although this is probably by chance).

The research on misinformation phenomena has matured greatly over the past two decades. Much work has been done in this area because the issues involved are so important. One of the most important developments is a consensus among researchers that no single factor is responsible for the effects. Whereas, in its infancy, researchers explained the misinformation effect as either a memory-based phenomenon or an artifact of the testing situation, most researchers today would agree that both of these factors—as well as others—are responsible for the observed effects. We believe that the activation-based perspective provides the most complete and succinct explanation that has been offered to date for understanding memory impairment effects in the misinformation effect literature. Its completeness in covering the issues and its generality in extending to other phenomena make it a useful framework for understanding issues relating to the misinformation effect specifically and to cognition generally. Future researchers in this area may find it useful to adopt this perspective for understanding the existing data and in designing future studies.

REFERENCES

- ANDERSON, J. R. (1974a). Retrieval of propositional information from long-term memory. *Cognitive Psychology*, **6**, 451-474.
- ANDERSON, J. R. (1974b). Verbatim and propositional representation of

- sentences in immediate and long-term memory. *Journal of Verbal Learning & Verbal Behavior*, **13**, 149-162.
- ANDERSON, J. R. (1976). *Language, memory, and thought*. Hillsdale, NJ: Erlbaum.
- ANDERSON, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- ANDERSON, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Erlbaum.
- ANDERSON, J. R., & BOWER, G. H. (1973). *Human associative memory*. Washington, DC: Winston.
- ANDERSON, J. R., & REDER, L. M. (in press). The fan effect: New results and new theories. *Journal of Experimental Psychology: General*.
- BARNES, J. M., & UNDERWOOD, B. J. (1959). "Fate" of first-list associations in transfer theory. *Journal of Experimental Psychology*, **58**, 97-105.
- BEKERIAN, D. A., & BOWERS, J. M. (1983). Eyewitness testimony: Were we misled? *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **9**, 139-145.
- BELLI, R. F. (1989). Influences of misleading postevent information: Misinformation interference and acceptance. *Journal of Experimental Psychology: General*, **118**, 72-85.
- BELLI, R. F. (1993). Failure of interpolated tests in inducing memory impairment with final modified tests: Evidence unfavorable to the blocking hypothesis. *American Journal of Psychology*, **106**, 407-427.
- BELLI, R. F., LINDSAY, D. S., GALES, M. S., & MCCARTHY, T. T. (1994). Memory impairment and source misattribution in postevent misinformation experiments with short retention intervals. *Memory & Cognition*, **22**, 40-54.
- BELLI, R. F., WINDSCHITL, P. D., MCCARTHY, T. T., & WINFREY, S. E. (1992). Detecting memory impairment with a modified test procedure: Manipulating retention interval with centrally presented event items. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **18**, 356-367.
- BOWER, G. H. (1996). Reactivating a reactivation theory of implicit memory. *Consciousness & Cognition*, **5**, 27-72.
- BOWERS, J. M., & BEKERIAN, D. A. (1984). When will postevent information distort eyewitness testimony? *Journal of Applied Psychology*, **69**, 466-472.
- BOWMAN, L. L., & ZARAGOZA, M. S. (1989). Similarity of encoding context does not influence resistance to memory impairment following misinformation. *American Journal of Psychology*, **102**, 249-264.
- CECI, S. J., & BRUCK, M. (1993). Suggestibility of the child witness: A historical review and synthesis. *Psychological Bulletin*, **113**, 403-439.
- CECI, S. J., ROSS, D. F., & TOGLIA, M. P. (1987). Suggestibility of children's memory: Psycholegal implications. *Journal of Experimental Psychology: General*, **116**, 38-49.
- CHANDLER, C. C. (1989). Specific retroactive interference in modified recognition tests: Evidence for an unknown cause of interference. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 256-265.
- CHANDLER, C. C. (1991). How memory for an event is influenced by related events: Interference in modified recognition tests. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **17**, 115-125.
- CHANDLER, C. C., & GARGANO, G. J. (1995). Item-specific interference caused by cue-dependent forgetting. *Memory & Cognition*, **23**, 701-708.
- COLLINS, A. M., & LOFTUS, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, **82**, 407-428.
- CROWDER, R. G. (1976). *Principles of learning and memory*. Hillsdale, NJ: Erlbaum.
- GARRY, M., MANNING, C. G., LOFTUS, E. F., & SHERMAN, S. J. (1996). Imagination inflation: Imagining a childhood event inflates confidence that it occurred. *Psychonomic Bulletin & Review*, **3**, 208-214.
- GREENBERG, R., & UNDERWOOD, B. J. (1950). Retention as a function of stage of practice. *Journal of Experimental Psychology*, **40**, 452-457.
- GREENE, E., FLYNN, M. S., & LOFTUS, E. F. (1982). Inducing resistance to misleading information. *Journal of Verbal Learning & Verbal Behavior*, **21**, 207-219.
- HYMAN, I. E., JR., HUSBAND, T. H., & BILLINGS, F. J. (1995). False memories of childhood experiences. *Applied Cognitive Psychology*, **9**, 181-197.
- JACOBY, L. L., WOLOSHYN, V., & KELLEY, C. M. (1989). Becoming famous without being recognized: Unconscious influences of memory produced by divided attention. *Journal of Experimental Psychology: General*, **118**, 72-85.

- JOHNSON, M. K. (1988). Reality monitoring: An experimental phenomenological approach. *Journal of Experimental Psychology: General*, **117**, 390-394.
- JOHNSON, M. K., HASHTRUDI, S., & LINDSAY, D. S. (1993). Source monitoring. *Psychological Bulletin*, **114**, 3-28.
- JOHNSON, M. K., & RAYE, C. L. (1981). Reality monitoring. *Psychological Review*, **88**, 67-85.
- KAMAS, E. N., & REDER, L. M. (1994). The role of familiarity in cognitive processing. In E. O'Brien & R. Lorch (Eds.), *Sources of coherence in text comprehension: A festschrift in honor of Jerome L. Myers* (pp. 177-202). Hillsdale, NJ: Erlbaum.
- KAMAS, E. N., REDER, L. M., & AYERS, M. S. (1996). Partial matching in the Moses illusion: Response bias not sensitivity. *Memory & Cognition*, **24**, 687-699.
- KELLEY, C. M., & JACOBY, L. L. (1996). Memory attributions: Remembering, knowing, and feeling of knowing. In L. M. Reder (Ed.), *Implicit memory and metacognition* (pp. 287-307). Hillsdale, NJ: Erlbaum.
- KEPPEL, G., & UNDERWOOD, B. J. (1962). Proactive inhibition in short-term retention of single items. *Journal of Verbal Learning & Verbal Behavior*, **1**, 153-161.
- LINDSAY, D. S. (1990). Misleading suggestions can impair eyewitnesses' ability to remember event details. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 1077-1083.
- LINDSAY, D. S. (1991). CHARMed but not convinced: Comment on Metcalfe (1990). *Journal of Experimental Psychology: General*, **120**, 101-105.
- LINDSAY, D. S., & JOHNSON, M. K. (1987). Reality monitoring and eyewitness suggestibility: Young children's ability to discriminate among memories from different sources. In J. Ceci, M. P. Toglia, & D. F. Ross (Eds.), *Children's eyewitness memory* (pp. 92-121). New York: Springer-Verlag.
- LINDSAY, D. S., & JOHNSON, M. K. (1989a). The eyewitness suggestibility effect and memory for source. *Memory & Cognition*, **17**, 349-358.
- LINDSAY, D. S., & JOHNSON, M. K. (1989b). The reversed eyewitness suggestibility effect. *Bulletin of the Psychonomic Society*, **27**, 111-113.
- LINDSAY, D. S., & READ, J. D. (1994). Psychotherapy and memories of childhood sexual abuse: A cognitive perspective. *Applied Cognitive Psychology*, **8**, 281-338.
- LOFTUS, E. F. (1975). Leading questions and the eyewitness report. *Cognitive Psychology*, **7**, 560-572.
- LOFTUS, E. F. (1979). Reactions to blatantly contradictory information. *Memory & Cognition*, **7**, 368-374.
- LOFTUS, E. F. (1991). Made in memory: Distortions in recollection after misleading information. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 27, pp. 187-215). San Diego: Academic Press.
- LOFTUS, E. F., DONDERS, K., HOFFMAN, H. G., & SCHOOLER, J. W. (1989). Creating new memories that are quickly accessed and confidently held. *Memory & Cognition*, **17**, 607-616.
- LOFTUS, E. F., & HOFFMAN, H. G. (1989). Misinformation and memory: The creation of new memories. *Journal of Experimental Psychology: General*, **118**, 100-104.
- LOFTUS, E. F., & KETCHAM, K. (1994). *The myth of repressed memory*. New York: St. Martin's Press.
- 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 & Memory*, **4**, 19-31.
- LOFTUS, E. F., & PALMER, J. C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning & Verbal Behavior*, **13**, 585-589.
- LOFTUS, E. F., SCHOOLER, J. W., & WAGENAAR, W. (1985). The fate of memory: Comment on McCloskey and Zaragoza. *Journal of Experimental Psychology: General*, **114**, 375-380.
- LOVETT, M. C., & ANDERSON, J. R. (1996). History of success and current context in problem solving: Combined influences on operator selection. *Cognitive Psychology*, **31**, 168-217.
- MCCLELLAND, J. L., & RUMELHART, D. E. (1986). *Parallel distributed processing: Explorations in the microstructures of cognition: Vol. 2. Psychological and biological models*. Cambridge, MA: MIT Press.
- MCCLOSKEY, M., & ZARAGOZA, M. (1985a). Misleading postevent information and memory for events: Arguments and evidence against memory impairment hypotheses. *Journal of Experimental Psychology: General*, **114**, 1-16.
- MCCLOSKEY, M., & ZARAGOZA, M. (1985b). Postevent information and memory: Reply to Loftus, Schooler, and Wagenaar. *Journal of Experimental Psychology: General*, **114**, 381-387.
- MCDERMOTT, K. B. (1996). The persistence of false memories in list recall. *Journal of Memory & Language*, **35**, 212-230.
- MCGEOCH, J. A. (1932). Forgetting and the law of disuse. *Psychological Review*, **39**, 352-370.
- MCGEOCH, J. A. (1942). *The psychology of human learning*. New York: Longmans, Green.
- MCGEOCH, J. A., & McDONALD, W. T. (1931). Meaningful relation and retroactive inhibition. *American Journal of Psychology*, **43**, 579-588.
- MELTON, A. W., & IRWIN, J. M. (1940). The influence of degree of interpolated learning on retroactive inhibition and the overt transfer of specific responses. *American Journal of Psychology*, **53**, 173-203.
- METCALFE, J. (1990). Composite holographic associative recall model (CHARM) and blended memories in eyewitness testimony. *Journal of Experimental Psychology: General*, **119**, 145-160.
- METCALFE, J. (1991). Representations, predictions and remembrances in CHARM: A reply to Lindsay (1991). *Journal of Experimental Psychology: General*, **120**, 313-315.
- METCALFE, J., & BJORK, R. A. (1991). Composite models never (well hardly ever) compromise: Reply to Schooler and Tanaka (1991). *Journal of Experimental Psychology: General*, **120**, 203-210.
- METCALFE, J., SCHWARTZ, B. L., & JOAQUIM, S. G. (1993). The cue familiarity heuristic in metacognition. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **19**, 851-861.
- MILLER, N. E., & STEVENSON, S. S. (1936). Agitated behavior of rats during experimental extinction and a curve of spontaneous recovery. *Journal of Comparative Psychology*, **21**, 205-231.
- MITCHELL, K. J., & ZARAGOZA, M. S. (1996). Repeated exposure to suggestion and false memory: The role of contextual variability. *Journal of Memory & Language*, **35**, 246-260.
- MORTON, J., HAMMERSLEY, R. H., & BEKERIAN, D. A. (1985). Headed records: A model for memory and its failures. *Cognition*, **20**, 1-23.
- NHOUYVANISVONG, A., & REDER, L. M. (in press). Rapid feeling-of-knowing: A strategy selection mechanism. In V. Y. Yzerbyt, G. Lories, & B. Dardenne, *Metacognition: Cognitive and social dimensions*. London: Sage.
- OSGOOD, C. E. (1949). The similarity paradox in human learning: A resolution. *Psychological Review*, **56**, 132-143.
- PAYNE, D. G., ELIE, C. J., BLACKWELL, J. M., & NEUSCHATZ, J. S. (1996). Memory illusions: Recalling, recognizing, and recollecting events that never occurred. *Journal of Memory & Language*, **35**, 261-285.
- PAYNE, D. G., TOGLIA, M. P., & ANASTASI, J. S. (1994). Recognition performance level and the magnitude of the misinformation effect in eyewitness memory. *Psychonomic Bulletin & Review*, **1**, 376-382.
- POSTMAN, L., & UNDERWOOD, B. J. (1973). Critical issues in interference theory. *Memory & Cognition*, **1**, 19-40.
- RANTZEN, A., & MARKHAM, R. (1992). The reversed eyewitness testimony design: More evidence for source monitoring. *Journal of General Psychology*, **119**, 37-43.
- RATCLIFF, R. A., & MCKOON, G. (1981). Does activation really spread? *Psychological Review*, **88**, 454-462.
- READ, J. D. (1996). From a passing thought to a false memory in 2 minutes: Confusing real and illusory events. *Psychonomic Bulletin & Review*, **3**, 105-111.
- REDER, L. M., & ANDERSON, J. R. (1980). A partial resolution of the paradox of interference: The role of integrating knowledge. *Cognitive Psychology*, **12**, 447-472.
- REDER, L. M., & GORDON, J. S. (1997). Subliminal perception: Nothing special, cognitively speaking. In J. Cohen & J. Schooler (Eds.), *Cognitive and neuropsychological approaches to the study of consciousness* (pp. 125-134). Mahwah, NJ: Erlbaum.
- REDER, L. M., & KUSBIT, G. W. (1991). Locus of the Moses illusion: Imperfect encoding, retrieval, or match? *Journal of Memory & Language*, **30**, 385-406.
- REDER, L. M., NHOUYVANISVONG, A., AYERS, M. S., SCHUNN, C. D., ANGSTADT, P., & HIRAKI, K. (1998). *Modeling the word frequency mirror effect in a remember/know paradigm*. Manuscript submitted for publication.
- REDER, L. M., NHOUYVANISVONG, A., SCHUNN, C. D., AYERS, M. S.,

- ANGSTADT, P., & HIRAKI, K. (1997). Modeling the mirror effect in a continuous remember/know paradigm. *Proceedings of the Nineteenth Annual Meeting of the Cognitive Science Society* (pp. 644-649). Mahwah, NJ: Erlbaum.
- REDER, L. M., & RITTER, F. E. (1992). What determines initial feeling of knowing? Familiarity with question terms, not with the answer. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **18**, 435-451.
- REDER, L. M., & ROSS, B. H. (1983). Integrated knowledge in different tasks: The role of retrieval strategy on fan effects. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **9**, 55-72.
- REDER, L. M., & SCHUNN, C. D. (1996). Metacognition does not imply awareness: Strategy choice is governed by implicit learning and memory. In L. M. Reder (Ed.), *Implicit memory and metacognition* (pp. 45-77). Hillsdale, NJ: Erlbaum.
- REDER, L. M., & WIBLE, C. (1984). Strategy use in question-answering: Memory strength and task constraints on fan effects. *Memory & Cognition*, **12**, 411-419.
- ROEDIGER, H. L., III, JACOBY, D., & McDERMOTT, K. B. (1996). Misinformation effects in recall: Creating false memories through repeated retrieval. *Journal of Memory & Language*, **35**, 300-318.
- ROEDIGER, H. L., III, & McDERMOTT, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 803-814.
- ROEDIGER, H. L., III, & SCHMIDT, S. R. (1980). Output interference in the recall of categorized and paired associate lists. *Journal of Experimental Psychology: Human Learning & Memory*, **6**, 91-105.
- SCHOOLER, J. W., FOSTER, R. A., & LOFTUS, E. F. (1988). Some deleterious consequences of the act of recollection. *Memory & Cognition*, **16**, 243-251.
- SCHOOLER, J. W., & TANAKA, J. W. (1991). Composites, compromises, and CHARM: What is the evidence for blend memory representations? *Journal of Experimental Psychology: General*, **120**, 96-100.
- SCHUNN, C. D., REDER, L. M., NHOUYVANISVONG, A., RICHARDS, D. R., & STROFFOLINO, P. J. (1997). To calculate or not calculate: A source activation confusion model of problem familiarity's role in strategy selection. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **23**, 3-29.
- SMITH, A. D. (1971). Output interference and organized recall from long-term memory. *Journal of Verbal Learning & Verbal Behavior*, **10**, 400-408.
- TULVING, E. (1985). Memory and consciousness. *Canadian Psychologist*, **26**, 1-12.
- TVERSKY, B., & TUCHIN, M. (1989). A reconciliation of the evidence on eyewitness testimony: Comments on McCloskey and Zaragoza. *Journal of Experimental Psychology: General*, **118**, 86-91.
- UNDERWOOD, B. J. (1957). Interference and forgetting. *Psychological Review*, **64**, 49-60.
- UNDERWOOD, B. J. (1965). False recognition produced by implicit verbal responses. *Journal of Experimental Psychology*, **70**, 122-129.
- WARREN, A. R., & LANE, P. (1995). Effects of timing and type of questioning on eyewitness accuracy and suggestibility. In M. S. Zaragoza, J. R. Graham, G. C. N. Hall, R. Hirschman, & Y. S. Ben-Porath (Eds.), *Memory and testimony in the child witness* (pp. 44-60). Thousand Oaks, CA: Sage.
- WEINGARDT, K. R., LOFTUS, E. F., & LINDSAY, D. S. (1995). Misinformation revisited: New evidence on the suggestibility of memory. *Memory & Cognition*, **23**, 72-82.
- WELLS, G. L., & TURTLE, J. W. (1987). Eyewitness testimony research: Current knowledge and emergent controversies. *Canadian Journal of Behavioral Science*, **19**, 363-388.
- WHEELER, M. A. (1995). Improvement in recall over time without repeated testing: Spontaneous recovery revisited. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 173-184.
- WINDSCHITL, P. D. (1996). Memory for faces: Evidence of retrieval-based impairment. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **22**, 1101-1122.
- ZARAGOZA, M. S., & KOSHMIDER, J. W., III (1989). Misled subjects may know more than their performance implies. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 246-255.
- ZARAGOZA, M. S., & LANE, S. M. (1994). Source misattributions and the suggestibility of eyewitness memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **20**, 934-945.
- ZARAGOZA, M. S., MCCLOSKEY, M., & JAMIS, M. (1987). Misleading postevent information and recall of the original event: Further evidence against the memory impairment hypothesis. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **13**, 36-44.
- ZARAGOZA, M. S., & MITCHELL, K. J. (1996). Repeated exposure to suggestion and the creation of false memories. *Psychological Science*, **7**, 294-300.

NOTES

1. Another common procedure is to have subjects read a narrative describing the events in the slides and to imbibe the misinformation within that narrative.

2. In principle, the source could be retrieved more directly, rather than inferred, if one made a conscious effort to encode the source at the time of presentation (e.g., repeating to oneself "The yield sign was in the slides"). This would presumably create a concept for this statement, which could then be retrieved at test. However, this would really just be a case of enriching the set of contextual associations to the encoding event.

3. This idea is analogous to the situation in which one edits a text file on a computer and then, without making a backup of the original, re-saves the file after the changes have been made. The new file is then the only record in existence, because the original has been permanently overwritten.

4. The encoding conditions for the interpolated faces were not the same across all experiments, but, within each experiment, they were held constant. In each experiment, impairment effects were found in some conditions but not in others, and it was always the case that impairment effects were found at delays of 45 min or less and were not found after 48 h or more, regardless of encoding conditions.

5. No other model we know of makes this assumption; it may or may not be unique.

6. Although one does not always find interference because of strategic effects (see, e.g., Reder & Anderson, 1980; Reder & Ross, 1983; Reder & Wible, 1984), the effects can always be understood and predicted within the same computational framework (see Anderson & Reder, in press, for more discussion of this issue).

7. ACT-R is a production system, whereas SAC only describes processing of declarative memory, so there is not a one-to-one mapping in assumptions.

(Manuscript received September 28, 1996;
revision accepted for publication September 9, 1997.)