COGNITIVE NEUROPSYCHOLOGY, 1999, 16 (3/4/5), 489–508

FALSE MEMORIES AND SOURCE MONITORING

Rahul M. Dodhia and Janet Metcalfe
Columbia University, New York, USA

This article investigates the relation between misremembering and source judgments in both the misleading information paradigm and the false memory paradigm. A computational model, CHARM (Composite Holographic Associative Recall Model), is used to simulate source monitoring in both paradigms. Despite the fact that CHARM stores memories in a composite memory trace, it is shown that the model can account for source judgements, and can explain the discrepancy between the source judgement and the recognition data in the misleading information paradigm. It also can account for the basic phenomena of the false memory paradigm, wherein thematically related items induce a memory for a nonpresented but prototypical critical item. In two experiments linking these two research lines, we presented the critical item in a different list from that used to induce the false memory effect. Although the model predicted that the presentation of the critical item should increase the false memory effect, its presentation inhibited false memories instead—but only with particular word lists and in certain treatment combinations. It seems likely that the presence of the critical item in an alternate list allowed people to use an exclusionary rule to inhibit the false memories. Such a rule would be straightforward to implement in CHARM, and could allow the model to account for this false memory suppression effect.

INTRODUCTION

Although misremembering may be produced in a number of different ways, there are two basic paradigms that have enjoyed extensive experimental investigation, and which are considered here. The first is the classic *misleading information paradigm*, in which misinformation is actually presented to the person, and it is this presented information that is the cause of the person's memory errors. The second is the *false memory paradigm*, in which the misremembered information is not presented. Instead, the false memories result from inferences or constructions that the person creates internally, and then later mistakes for events that happened externally. In this article, we will address both of these misremembering paradigms. It has been shown, in the classic misleading information paradigm, that if people are asked to make judgements concerning the origin or source (Johnson, Hashtroudi, & Lindsay, 1993) of the would-be misleading information they are sometimes able to sharpen their focus, and discount the misleading information by attributing it to its correct source. Thus, in at least one of these paradigms, source judgements can be used to enhance people's memory, and eliminate false information (Lindsay & Johnson, 1989; Zaragoza & Koshmider, 1989). We investigate whether such judgements might also be efficacious in the other paradigm.

Requests for reprints should be addressed to Janet Metcalfe, Department of Psychology, 401 Schermerhorn Hall, Columbia University, New York, 10027, USA (E-mail: metcalfe@paradox.psych.columbia.edu).

This research was supported by National Institute of Mental Health Grant MH 48066 to Janet Metcalfe. We thank Lisa Schwartz for her help with Experiment 2. We also thank Walter Mischel, Kenneth A. Norman, Henry L. Roediger III, and Daniel L. Schacter for their helpful comments on an earlier version of the manuscript.

In the misleading information paradigm (of which there are a number of variants) an incident, such as a crime, is witnessed and then later, if he or she is in the experimental condition, the person is given some misleading information about some detail in the incident. For example, if the person originally saw a car go through a red light, it might later be suggested that the vehicle had been a van. When given a two-alternative forced choice test including the correct object and the misleading suggestion, people who received the misleading information performed worse than did people who received no misleading information (Loftus, 1979). The CHARM model has been used, in the past, to model situations such as this (Metcalfe, 1990). It explained these, and related, findings in terms of the implications of retrieving a superimposed representation that contains both the original event (the car) superimposed on the misleading suggestion (the van). The effect is analogous to a TV screen showing simultaneously two different programmes, or to a photograph that has been double-exposed. When those two exact alternatives are retrieved and contrasted in this superimposed manner, at the time of test, the person has difficulty distinguishing them, and so performance suffers—as has been shown empirically.

A number of researchers have criticised and modified the original version of the misleading information paradigm. When the lures given at the time of test are not identical to the misleading item that was first suggested, people who received the misleading suggestion usually performed at about the same level as did control participants (e.g. McCloskey & Zaragoza, 1985; cf. Payne, Toglia, & Anastasi, 1994). The CHARM model—employing a composite memory trace—can predict these data (Metcalfe, 1990). Only the misleading item is superimposed on the correct item at time of retrieval, so other items tend not to interfere with the correct item (unless, of course, they are sufficiently similar to the original item, the suggested item, or both, so that confusion occurs because of a discriminability problem). There are a number of other experimental variants on this basic misleading information paradigm (see, for example, Belli, 1989; Ceci, Ross, & Toglia, 1987; Chandler, 1989; Christiaansen & Ochalek, 1983; Kroll & Timourian, 1986; Loftus & Hoffman, 1989; Tversky & Tuchin, 1989) to which the CHARM model has been applied with some success. The composite retrieval characteristic of the model has allowed it to account for all of the major recognition and recall findings reported in this paradigm, except for one.

The only major experimental finding within the classic misleading information paradigm to which the model has not been applied, as pointed out by Lindsay (1991), is source monitoring. Rather than being tested for only for simple recognition or recall, in the source monitoring paradigm, people are asked to remember and report the source of the memory probe item. This paradigm is of especial interest for three reasons. First, judgements of source are a key component to people's episodic memory. Second, these judgements allow people to clarify and improve their memory performance, and greatly offset their susceptibility to misleading information (Lindsay & Johnson, 1989). The third reason for our interest is that the manner in which the CHARM model retrieves a representation with information from multiple sources superimposed, would seem, at first blush, to disallow it from making source judgements. Much of the explanatory power of the CHARM model stems from the construct of a composite trace—including the basic explanation for the misleading information effect itself. The challenge to CHARM, then, is that it would seem that a model that uses a composite trace as a core construct (as opposed, say, to the more traditional separate-trace models) might seem intuitively to be unable to distinguish the sources. In the first section of this article, we will show, though, that CHARM can handle the basic source-monitoring results without difficulty.

In the second major section of this article we turn to the other misremembering paradigm. Roediger and McDermott (1995, and see Deese, 1959) have investigated a paradigm for creating false memories, in which the false memories are not explicitly implanted by an outside source, as in the classic paradigm; instead, the participant appears to create them by some internal process. Participants are presented with a list of semantic associates of a critical item. They do not study the critical item itself but, nevertheless, exhibit false memory for it at test. For example, a person might be presented with a list of semantically associated words such as *hill, valley, climb,* etc., and then be tested on a critical item such as *mountain*. This paradigm, which has been of great interest recently, will be described in greater detail later in the paper.

In the classic misleading information paradigm people can use source information to differentiate suggested from true memories about a particular scene (Lindsay & Johnson, 1989; Zaragoza & Koshmider, 1989). We conjectured that source judgements might also be useful in the false memory paradigm. To investigate this possibility, our first step was to model the effect of source-monitoring in the misleading information paradigm within the CHARM (composite, holographic, associative, recall, and recognition memory) model (Metcalfe Eich, 1982, 1985; Metcalfe, 1990, 1991a,b, 1993a,b, 1997; Metcalfe, Cottrell, & Mencl, 1992). We then showed that this simple model is able to account for the basic phenomena within the false memory paradigm. Finally, we used CHARM's source monitoring predictions to investigate the impact of the presence of source information in the false memory paradigm.

MODELLING SOURCE JUDGEMENTS WITH *CHARM* IN THE MISLEADING INFORMATION PARADIGM

Lindsay and Johnson (1989) conducted an experiment in the misleading information tradition, in which they showed that memory could be ameliorated by asking participants for source judgments. In their experiment, the participants viewed an initial event—a picture taken from a magazine. They then read a text that was a description of the picture but which included some misleading information (as well as other information). Then participants were either asked to make old/new recognition judgements about whether each of the probe items was in the original picture they had seen (in a manner similar to that used in the classic paradigm), or they were given a source judgement task. In the latter task, they were asked to say whether a given item was seen only in the picture, was only in the text, was in both the picture and the text, or was entirely new. To equate between the recognition judgement and the source judgement conditions, the probability of saying that an item was seen in the picture (as given directly in the recognition judgement case) was taken, in the source judgement condition, to be the sum of the probabilities of saying that it was in the picture only and that it was in both.

Technically the two tests were asking the same question: Was the probe in the picture? Nevertheless, the results were quite different. Misled participants given the recognition test were much more likely to false alarm to the misleading items that were presented only in the text than were their matched controls who were given no misleading suggestions. In the source judgement condition, though, misled participants were perfectly well able to say that these items were only in the text, and they did not exhibit a tendency to claim that they had been in the picture. These results, of course, are of practical interest. But they raise an interesting theoretical question as well: Can a model that produces a superimposed retrieved item account for this discriminative ability? And if it can, how does it do it?

In this section we will first demonstrate that there are conditions, mapping onto the experimental situation in a straightforward manner, that allow the CHARM model to produce both the recognition results and the source monitoring results. Not surprisingly, the model is unable to produce both patterns of data—those for recognition and for source judgement—if it does exactly the same thing. So examination of exactly how the model produced the two different sets of data is of interest.

Source Simulations of CHARM

Items are represented in CHARM as vectors, and the elements of these vectors correspond to the features of the items. The elements are sampled from a standard normal distribution and the vectors are usually normalised to make mathematical analysis easier. The vectors created in this way are independent of each other, so the similarity of one random vector to another, as measured by their dot product, is 0. The similarity of one vector to another can be manipulated, however, by setting their dot product to some fraction greater than 0. Items are associated by the operation of convolution, and the resulting vector is superimposed onto an existing trace, which is initialised to be a noisy random vector. Items can be retrieved from the trace by correlating it with a probe vector. The resulting vector is a combination of all the items with which the probe had been associated. This retrieved vector is then sent off to a decision routine that distinguishes among yes/no recognition, forced-choice recognition, recall, and so on, as dictated by the constraints of the situation. In yes/no recognition, the retrieved vector is matched to the probe itself. If the dot product is greater than a preset criterion, the probe is considered as having been recognised. In multiple-choice recognition, the retrieved item is matched to each of the presented alternatives, and the one that matches best is given as the selection chosen, as long as it exceeds a criterion (if, in the experiment, the option to not answer is given). In recall, the retrieved vector is matched to each of the vectors in the lexicon. The vector that is recalled is the vector that provides the largest dot product with the retrieved vector, if the dot product is above a preset criterion.

In our first set of simulations, every item was assumed to have source information linked to it, via the associative operation of convolution. Source information consists of the perceptual and other sensory details that are the setting in which an event occurs. For modelling purposes, **Source**, **S**, is abstracted into "features" just as content information is. For the purposes of demonstrating the detailed predictions of the model, context or source, and content will be represented in the following way: If someone sees a picture of a car, this will be considered to be encoded in the trace **Trace**, as:

$$\mathbf{Trace} = \mathbf{Car{*}Car} + \mathbf{Car{*}Source}_{\text{Picture}},$$

where **Car** is a vector representing the mental representation of the car, $\mathbf{Source}_{\text{Picture}}$ is a vector giving the mental representation of the picture source, and "*" is the operation for convolution. When the cue **Car** is given, retrieval occurs by correlating (#) the **Trace** with **Car**. In general, the result of retrieval by any cue, **Q**, which is correlated with an association between any two vectors such as **X*Y**, is that **Y** is retrieved to the extent that **Q** is similar to **X**, and **X** is retrieved to the extent that **Q** is similar to **Y**, and the two components are added to one another. This occurs for all associations in the composite memory trace (and all retrieved components are added together). Thus, in the present case, the result of correlating the cue **Car** with the **Trace** is a vector which contains both **Car** and $\mathbf{Source}_{\text{Picture}}$, specifically:

$$\mathbf{Car\#Trace} = \mathbf{Car\#[(Car{*}Car) + (Car{*}Source}_{\text{Picture}}\mathbf{]} = \mathbf{2Car} + \mathbf{Source}_{\text{Picture}} \quad (1)$$

Since the source of the car memory is returned in the retrieved vector, this vector can be used to make a source judgment. The source of an item is judged to be retrieved if the dot product of the source vector in question and the retrieved vector is above a set criterion.

In Lindsay and Johnson's (1989) experiments the first factor was type of study group—Control or Misled. The second factor was type of test—Recognition or Source Judgement. In phase 1, both groups looked at a slide containing various items, thus encoding them in a pictorial context. In phase 2, both groups read a narrative describing the slide, encoding items in a textual context. The Control group read a narrative that was completely faithful to the slide, but the Misled group read a narrative that had additional information in it, items that had not appeared in the slide. In phase 3, participants from both groups received either a yes/no recognition test or a source judgement test. The test items consisted of items that had appeared only in the slide (Picture Only), items that had appeared in both the slide and the narrative (Picture and Text), items that had appeared in the misleading narrative only (Text Only) and items that had neither appeared in the slide nor the narrative (New).

Method

To model this situation, a lexicon of 10 vectors of 63 elements each was created. Vectors 1 and 2 were designated **W** and were associated with the picture source $\mathbf{S}_p$ only. Vectors 3 and 4 were designated **X**, and were associated with both the picture source $\mathbf{S}_p$ and the text source $\mathbf{S}_t$, vectors 5 and 6 were designated to be **Y** and were associated with the text source $\mathbf{S}_t$ only, vectors 7 and 8 were designated **Z** and reserved to be new. Two additional items were added into the composite trace as noise, but were not tested. Four sets of this lexicon were created, each with different random vectors. To avoid repetition and enhance clarity, the traces below are shown with only one set of vectors. In the simulation all four sets were used. Two other vectors were also created—they were source vectors, specifically Picture Source $\mathbf{S}_p$ and Text Source $\mathbf{S}_t$. These were simply set up as independent random vectors.

Two traces were constructed, one for each of the study groups. The Control trace ($\mathbf{T}_c$) was constructed as follows:

$$\mathbf{T}_c = \mathbf{W}*\mathbf{W} + \mathbf{W}*\mathbf{S}_p + \mathbf{X}*\mathbf{X} + \mathbf{X}*\mathbf{S}_p + \mathbf{X}*\mathbf{X} + \mathbf{X}*\mathbf{S}_t + \text{noise}$$

The Misled trace ($\mathbf{T}_m$) was constructed as follows:

$$\mathbf{T}_m = \mathbf{W}*\mathbf{W} + \mathbf{W}*\mathbf{S}_p + \mathbf{X}*\mathbf{X} + \mathbf{X}*\mathbf{S}_p + \mathbf{X}*\mathbf{X} + \mathbf{X}*\mathbf{S}_t + \mathbf{Y}*\mathbf{Y} + \mathbf{Y}*\mathbf{S}_t + \text{noise}$$

The test items **W**, **X**, **Y**, and ***Z*** were each correlated with the traces. An item was classified as having been recognised if its similarity (measured by a dot product) to the vector it retrieved from the trace was above a criterion of 1.2. To enact the source judgement the probe was correlated with the trace and the vector that was retrieved was matched with the picture context and with the text context. If the dot product of either match was above 0.8, then the model recognised the probe as having occurred in the corresponding context. If both matches were above criterion, then the probe was recognised as having appeared twice, once in each context. If neither match was above criterion, then the probe was assumed to be new. The simulation was replicated 500 times, with different random vectors each time.

Results and Discussion

As can be seen from Fig. 1, the model, like people, produced a high rate of false "attribution to picture," in the recognition task only, to those items that were presented in the text only and were not given in the picture. In the source judgements task, however, this error was attenuated, both in people's judgements and in the model.

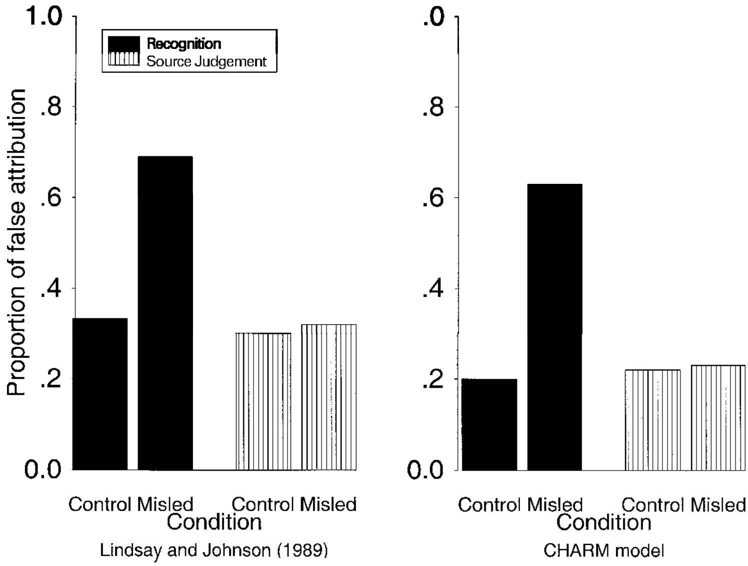

Fig. 1. *A comparison between Lindsay and Johnson (1989) and CHARM of the mean proportion of false attribution to "Picture" of "Text Only" items.*

The reason the model did this was that the recognition decision was based on matching the item retrieved to the probe only (as has been the method of recognition used in many past simulations with CHARM, e.g. Metcalfe, 1991a), and it did not consider the source. When only the recognition of the probe itself is considered, and the source is disregarded, a high false alarm rate to items presented in the non/pictorial context arises. In contrast, under the source judgement instructions, the model was specified to match the retrieved item to both the picture and the text source, and to make attribution decisions based on the extent of those matches. If the retrieved vector failed to match the picture source it would not be attributed to the picture. Because the source retrieved by the model for the misleading item probes was the text and not the picture (because it had never been associated with the picture, but had been associated with the text), the model did not have a high rate of misattribution of these probes to the picture. It is notable that, in both the recognition and the source judgment conditions, the item that was retrieved by the model was actually the same and the source information was present. In the recognition condition it was not considered in the decision, however, whereas it was in the source judgement condition.

Figure 2 shows the pattern of data concerned with attribution to the picture source, in the experiment and the model, for all probe items rather than only for the critical misleading items. The match of the model to the experimental data was good. Figure 3 shows the distribution of attribution to "new," both in the model and in the experiment.

Increased Emphasis on Source in Recognition Instructions

Lindsay and Johnson (1989) conducted a second experiment, related to the first, in which their concern was that people had not understood the recognition instructions. Perhaps the participants may have just thought that they were to call the items old if they had occurred at all in the experiment. Thus, in the follow-up experiment they emphasised that people were to say old *only* if the item occurred in the picture, and not otherwise. When Lindsay and Johnson clarified and emphasised that occurrence in the picture was crucial, they found that this manipulation had an effect in decreasing the false alarms to the items that were presented in the text but not in the picture context. In CHARM, by allowing the source assessment procedure to come into play in the recognition condition, we were easily able to simulate Lindsay and Johnson's follow-up experiment. As we will argue, when we discuss our own experiments, participants may sometimes use only a simple recognition process (which some researchers call a "familiarity" judgement, though this too involves retrieval in CHARM) when the context in which the probe was experienced is requested, but they may also, under certain circumstances, voluntarily use the source monitoring procedure to check their responses.

Source Judgements When the Misleading Items are Related to Original Items

An experiment related to that of Lindsay and Johnson (1989) was also conducted by Zaragoza and Koshmider (1989); in it, participants were presented with a series of slides followed by a narrative describing them. Unlike the Lindsay and Johnson experiments, though, in which the misleading items were unrelated additions, the misleading items in Zaragoza and Koshmider's experiment were highly related to particular original items. For example, a screwdriver in the slides would be replaced with a hammer in the narrative. Participants then received a source monitoring test in which the test items were presented as slides, rather than as text. Test probes could be either original test items or misleading test items. Participants had to decide whether they had seen each test probe in the picture, in the text only, or, if they could not remember the source, then whether the item was consistent or inconsistent with what they had seen in the experiment.

We conducted simulations of this paradigm in CHARM. The main difference between the previous simulations and these ones was that the misleading items, which were themselves autoconvolved and added into the composite memory

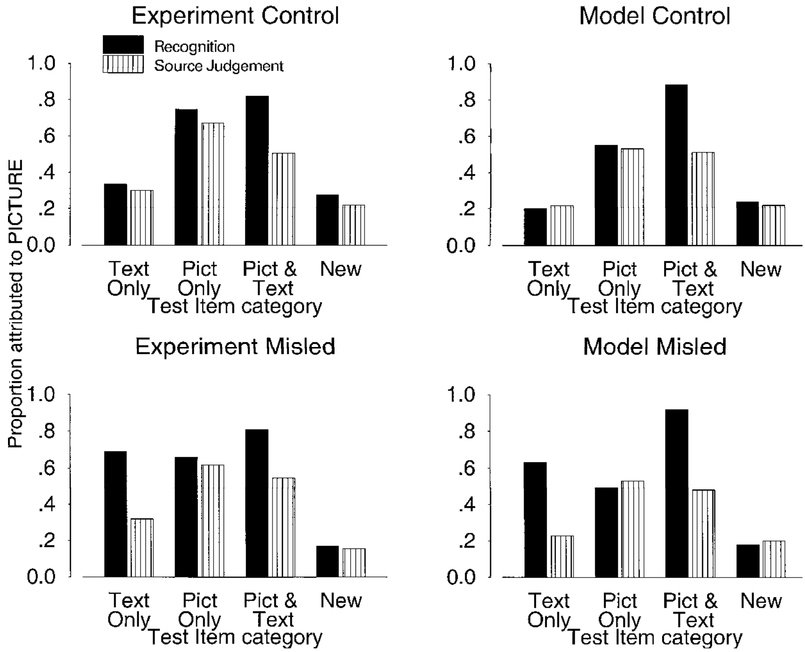

Fig. 2. *A comparison between Lindsay and Johnson (1989) and CHARM of the mean proportion of false attribution to "Picture" of all test items.*

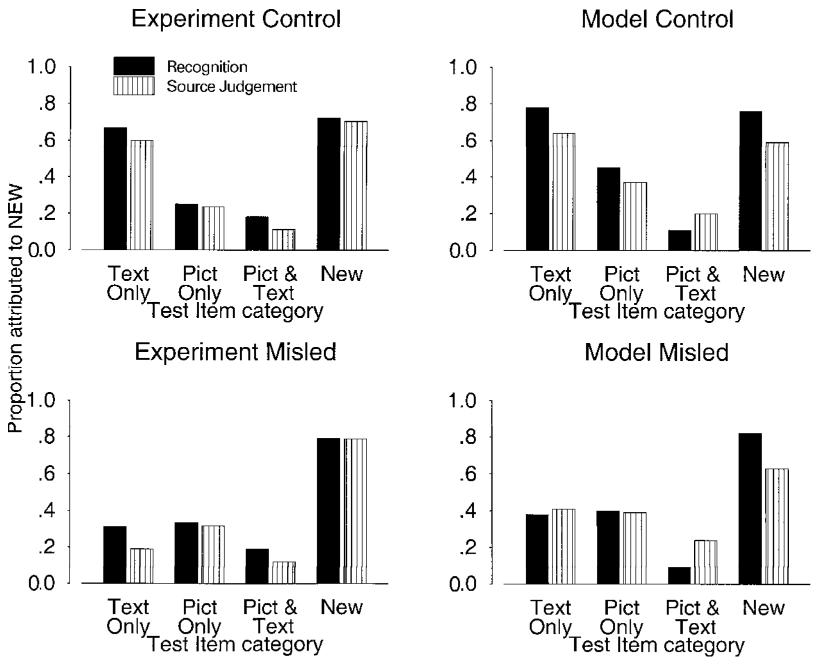

Fig. 3. *A comparison between Lindsay and Johnson's (1989) experiment and CHARM predictions of the mean proportion of false attribution to "New" of all test items.*

trace, and which were also convolved with a vector representing the narrative context, in this case were vectors that had considerable feature overlap with the vectors entered into the composite memory trace in association with the original context. The results of these simulations will not be reported in detail, except to say that they produced results that were consistent with the data reported by Zaragoza and Koshmider. Thus, regardless of whether the misleading items are related or unrelated to the original items, source judgements are useful in improving performance. The simple mechanisms in the CHARM model, used to account for these differences between recognition and source judgements, are sufficient to account for the results found experimentally.

The logic of incorporating source information into memory encoding and decisions, in at least one of these three forms, can also be implemented other models, such as those of Hintzman (1988), Ayers and Reder (1998), or Reyna and Brainerd (1995). These models have very different conceptual bases from CHARM, but can plausibly handle these source monitoring data. It is especially pleasing, though, that models relying on a composite memory trace can account for the source monitoring results since the nature of the composite trace and composite retrieval—which have many interesting conceptual spin-offs—makes this capacity unintuitive for this class of models.

THE FALSE MEMORY PARADIGM

The second half of this article focuses on the second paradigm for misremembering—the false memory paradigm—in which people are presented with a variety of related items, and tend to false alarm on critical items that were not presented, but which are highly related to those items that were. For example, a person might be presented with a list of items including *sugar, cake, honey,* and *chocolate*. Later they would be tested with the critical (unpresented) item *sweet*, which is a high associate of the presented items. Roediger and McDermott (1995) reported false alarm rates in recognition ranging from 58% to 81% for such critical items, though Roediger, Gallo, Watson, and Balota (1998) have recently shown much more variability. Robinson and Roediger (1997) showed that the rate of false attributions increases monotonically with increases in the number of associates in the study list. Norman and Schacter (1997) showed that older adults are more vulnerable to these kind of false memories than are younger adults.

Perhaps surprisingly, given the data with older adults, Korsakoff and non-Korsakoff amnesics showed reduced levels of false memories (Schacter, Verfaellie, Anes, & Racine, 1998). Interestingly enough, as the level of remembrance for true memories increased, the Korsakoff amnesics also showed an increase in their level of false memories. In the same experiment, though, with normal controls, as the level of true memory increased, the rate of false alarming to the critical item decreased.

Several experiments have shown that repetition of the study list (with normals) results both in an increase in true memory and a decrease in false memories (McDermott, 1996). Furthermore, other manipulations that increase the distinctiveness of the items in the to-be-remembered list also decrease the probability of false memories. Israel and Schacter (1997) showed that the pictorial encoding of the inducing list items—which were presented in the form of line drawings—decreased the false memory effect. However, there is some indication that neither repetition nor pictorial encoding are effective in obviating the effect in elderly subjects (Kensinger & Schacter, this issue; Koutstaal & Schacter, 1997; Schacter, Kaszniak, Kihlstrom, & Valdiserri, 1991).

CHARM and False Memories

The most straightforward way to handle the false memory paradigm in CHARM is to consider it to be a special case of prototype formation, which has been modelled formally in CHARM (Metcalfe Eich, 1982; and see Metcalfe, 1991b). In that paradigm, as in the false memory paradigm, a number of items that are highly similar to one another—call them $\mathbf{I}_1$, $\mathbf{I}_2$, $\mathbf{I}_3$, $\mathbf{I}_4$, etc., and to an item that might be called either a "critical item" in the false memory paradigm) or a prototype—call it $\mathbf{I}_p$ (for prototype)—are autoassociated and added into the composite trace. Thus, the composite trace is:

$$\mathbf{T} = \mathbf{I}_1 * \mathbf{I}_1 + \mathbf{I}_2 * \mathbf{I}_2 + \mathbf{I}_3 * \mathbf{I}_3 + \mathbf{I}_4 * \mathbf{I}_4 + \ldots + \text{noise} \quad (2)$$

The predictions of the model, in this situation, are straightforward. If $\mathbf{I}_p$ is used as the retrieval cue, each of the encoded **I***'s* will be retrieved as a function of their similarity to $\mathbf{I}_p$. Furthermore, they will all be superimposed (see Metcalfe, 1991b, for some illustrations of superimposed representations going back to Galton's (1878) seminal work on composite photography and "family resemblance"). $\mathbf{I}_p$, then, will retrieve something that closely resembles a prototype, and if that resultant retrieved item is matched to $\mathbf{I}_p$, itself, the extent of match can be very high indeed, depending upon the similarity of the encoded items ($\mathbf{I}_1$, $\mathbf{I}_2$, etc.) to $\mathbf{I}_p$, and also upon how many such items were encoded. If very few related items were entered into the trace, then the generalisation to the critical item will be low, but as more and more highly/related items are added into the composite trace the generalisation to the critical item increases. This finding in the model (see Metcalfe Eich, 1982, for simulations) corresponds nicely to Robinson and Roediger's (1997) findings.

It also follows that if people are remembering very few of the presented items, they may have encoded only a few items. Even though a large number of the related items might have been presented, if only a few of those items were encoded and entered into the trace, then the situation is the same as if only a few related items were presented. In the model, when one item is encoded in the list, and an item that is similar to it is presented as the cue, the retrieved item will match the cue item as a function of the degree of similarity between the two, but there is no convergent generalisation from multiple items. Similarly, with only a few encoded items, there may be insufficient information in the composite trace to support the generalisation to the critical items, and no false memory effect will be manifest. Thus some degree of true memory is needed if people are to demonstrate false memories, since the "false" memories themselves depend upon superimposed retrieval of the true memories. As the level of overall memory increased with the amnesics in Schacter et al.'s (1998) study, presumably the number of encoded items increased, allowing generalisation to the false memory. Thus, the finding that amnesics may fail to show a false memory effect with low levels of true memory, but then show the false memory effect with higher levels of true memory, may be strongly related to the finding of Robinson and Roediger (1997)—that false memories increase with an increase in the number of associates presented. Furthermore, as the inducing items are made more and more similar to one another and to the critical item, the model predicts that false memory for the critical item should also increase.

At the other extreme, if the individual items in the list are extremely well discriminated, that is, they are either not highly related to one another, or if they are encoded such that their distinctive features are emphasised, the false memory effect should be offset. This can be captured in the model either by using items that are not highly similar to one another (see Metcalfe Eich, 1982, for simulations), or by allowing a high proportion of distinctive features (and not just the most typical features) to be encoded by a selective feature-sampling process (see Metcalfe, 1991b, for details of feature sampling in the model, and some of the implications). High levels of distinctive encoding could allow discrimination of the critical item from the presented items. Thus, at extremely high levels of learning of the individual items—as long as that learning entailed emphasis on the distinctive features—the false memory effect may disappear. (There are conditions under which the false memories may not disappear, however. The explanation forwarded here is that, with many presentations, people are intentionally forging distinctive mental representations that distinguish the presented items from one another and from the critical item. Looking for, and encoding, the differences among the items is likely to be under voluntary control, however, and if the participants in the experiment are not doing it, but are simply strengthening the canonical representational forms of the presented items, then high levels of true memory may still give rise to high levels of false memory, as in the experiments conducted on ageing people, cited earlier.)

This interpretation is consistent with that of Schacter, Israel, and Racine (1999).

There are other experimental manipulations that appear to give rise to distinctive encodings, where distinctive means that the features that are different across items are particularly emphasised. Pictures may evoke a more distinctive and unique event encoding than do words. Several experiments have shown that pictorial encoding is less likely to evoke false memories than is verbal encoding (Israel & Schacter, 1997; Schacter et al., 1999). The idea that the difference between pictures and words, in this paradigm, is a function of the distinctiveness of the two kinds of materials, leaves open the possibility that one might be able to choose pictures to be nondistinctive, in which case they may be as good as words at producing false memories. The disappearance of the false memory effects at high levels of memory, then, is for reasons quite different from those accounting for its failure to occur at very low levels of memory.

Theoretical Predictions about False Memories and Source

There may be other situations that could allow a person to discriminate a false memory from events that were present in the episode under study. Our hypothesis, in the present experiment, was that source information may provide a clue that, at least under some circumstances, may serve as discriminative information.

The paradigm investigated here was devised to produce potentially paradoxical effects. We thought that by presenting the critical item itself we might, under some very constrained situations, set the stage for people to inhibit false memories concerning that item. However, although these were our expectations, they were *not* the straightforward predictions of CHARM, and so this paradigm holds especial interest for us. The paradigm was one in which the human participants (and the model) were presented with two lists. One list contained the highly related items like those usually used to induce a false memory effect. We will call this list the Inducing List. In the Other List we either presented (in the Experimental Condition) or did not present (in the Control Condition) the critical item.

To instantiate this situation in the model, each of the related items in the Inducing List is autoconvolved, and also convolved with the Inducing List source vector, and all of these convolutions are added into the composite memory trace. Thus, the situation resembles that given in Equation 2, except that each item is also convolved with its list source vector. The items in the Other List are also autoconvolved and each is convolved with an item representing the Other List source. In the equation given below, though (because, aside from the critical item itself, the content of the vectors in the other list is irrelevant) we will simply call these associations that comprise the associations from the Other List "noise". The Inducing List Source is notated $\mathbf{S}_i$ (for inducing); the source vector representing the Other List is notated $\mathbf{S}_o$ (for other). In the Experimental Condition the critical item, which is notated as $\mathbf{I}_p$ (or prototype item, as given in Equation 2) is encoded as a member of the Other List. In the control condition the critical item is not encoded at all. Schematically, then, the experimental trace, $\mathbf{T}_e$, and the control trace, $\mathbf{T}_c$, are as follows:

$$\mathbf{T}_e = \mathbf{I}_1{*}\mathbf{I}_1 + \mathbf{I}_1{*}\mathbf{S}_i + \mathbf{I}_2{*}\mathbf{I}_2 + \mathbf{I}_2{*}\mathbf{S}_i + \mathbf{I}_3{*}\mathbf{I}_3 + \mathbf{I}_3{*}\mathbf{S}_i + \ldots + \text{noise (from the rest of the Inducing List and from preexperimental sources)} \ldots + \mathbf{I}_p{*}\mathbf{I}_p + \mathbf{I}_p{*}\mathbf{S}_o + \text{noise (from Other List)}$$

$$\mathbf{T}_c = \mathbf{I}_1{*}\mathbf{I}_1 + \mathbf{I}_1{*}\mathbf{S}_i = \mathbf{I}_2{*}\mathbf{I}_2 + \mathbf{I}_2{*}\mathbf{S}_i + \mathbf{I}_3{*}\mathbf{I}_3 + \mathbf{I}_3{*}\mathbf{S}_i + \ldots + \text{noise (from the rest of the Inducing List and from preexperimental sources)} \ldots + \text{noise (from Other List)}$$

The crucial difference between the two traces is in whether or not the critical item itself ($\mathbf{I}_p{*}\mathbf{I}_p$), as well as its association to the Other List source vector ($\mathbf{I}_p{*}\mathbf{S}_o$), is or is not included

At time of retrieval, the critical item, $\mathbf{I}_p$ is correlated with the composite trace, and, under recognition instructions, the vector that is thereby retrieved is compared to the probe (as is usual in recognition in the model). The match between the retrieved item and the critical item probe, in the Control Condition, is high (by virtue of all the highly related items retrieved from the Inducing List). But the match between the retrieved item and the critical item probe, in the Experimental Condition, is even higher, because in addition to all of the information retrieved in the Control Condition, the presented critical item itself is also retrieved. Thus CHARM—like all other models of simple recognition—predicts that under conditions of recognition—as long as people are only considering the information about the content of the item itself and are disregarding source information, the probability of saying "yes" to the critical item should be higher in the experimental condition than in the control condition.

What about the predictions when the task is source monitoring? The answer is straightforward if people are making a judgement as given in the simulations presented in the first section of this article. In the Control Condition, there is a great deal of retrieved information converging on the Inducing List as the source of the critical item, since each of the inducing items was associated with the Inducing List Source vector. There is no information pointing to any other source for the critical item. Thus, the model predicts that people should think that the critical item came from the Inducing List, and not from the Other List or from both lists. The model will sometimes allow that the source is "neither" if the resonance criterion is not reached.

In contrast, in the Experimental Condition, while all of the information pointing to the Inducing List as the source is present in the Experimental Condition as it was in the Control Condition, in addition, the presented critical item was associated with the Other List source, and hence, via this association the probe also retrieves the Other List source. Given that both of these sources are retrieved, then, the model predicts that people should give all three responses—Inducing List, Other List, and especially BOTH lists—as the source of the critical item.

Thus, to summarise, the predictions of the model were that under recognition conditions, people should be more likely to affirm that the critical item was in the Inducing List if it was presented in another list. The experimental condition should produce higher false recognition of the critical item than should the Control Condition. If asked for the source of the critical item in the Control Condition, people should say it was from the Inducing List. In the Experimental Condition they should say that it was present in both the Inducing List and the Other List.

Experiments 1 and 2

We ran two separate experiments, both of which conform to the basic design such that the critical item is presented in the alternative list in the Experimental Condition, or is not presented in the Control Condition. The results of these experiments will be presented individually because certain interactions appear in one experiment and not in the other. The first experiment was run on the computer, using PsyScope; the second was run on flash cards. The second experiment was conducted because the first experiment did not allow us to investigate the fate of the exact same materials in the experimental and control conditions. Roediger et al. (1998) have shown that there can be extreme differences in results in the false memory paradigm as a function of the exact word lists used, and so we were concerned (correctly, as it turns out) that this aspect of the experiment be carefully controlled before reaching any conclusions. In the first experiment participants were tested both for source judgements and for yes/no recognition. In the second experiment we tested only for yes/no recognition.

Methods

Participants. Participants were students in an Introductory Psychology course at Columbia University, who received partial course credit for participating in this experiment, or were paid volunteers (also Columbia University students) who were given $10 for their efforts. There were 67 participants in Experiment 1 and 26 in Experiment 2. Data from two participants, in Experiment 1, were eliminated because these individuals did not follow the instructions.

Design. Participants studied two lists in every trial. Each list consisted of seven associated "inducing" words that were related to a critical item and two words that were unrelated to the general theme of the list. In the Experimental Conditions, one of these two words was the critical word for the alternative list. In the Control Conditions, both of these words were unrelated both to the theme of the list in which they were presented and to the theme of the alternative list. The two lists were not related to each other. Thus, each participant was tested on what was basically a 2 × 2 within-subjects design. The two factors were: Presentation of Critical Item (presented in the alternate list in the Experimental Condition or not presented in the Control Condition) and Test List (either List 1 or List 2).

Test List 1 means that the inducing words—*cake, sugar, honey,* etc.—were presented in List 1, and that subjects were instructed, in recognition, to say yes only to words that had occurred in List 1. On experimental trials, the critical word "*sweet*" would be presented in the second list; on control trials nothing else related to the first list would be presented in the second list. Conversely, Test List 2 implies that List 2 was the inducing list and that the critical item appeared—if the particular trial in question was one in the Experimental Condition—in the first list.

In Experiment 1, 33 participants were tested exclusively on a yes/no recognition task and 32 participants received exclusively a 4-choice source judgement task. In Experiment 2, all participants were tested on yes/no recognition only.

There were several other factors included, for control purposes, in the first experiment. A between-participants factor was Trial Type. For 33 participants, the list pairs comprising each trial were structured such that on each trial only one critical item was presented. The other 32 participants saw trials in which both lists contained critical items, or neither contained critical items. These two methods of linking the two lists were initially conducted as separate mini-experiments and run sequentially during the semester, but they are presented here as one experiment since there were no effects of this variable. Another factor was the Set of Materials used. Set B reversed the ordering of pairs of linked lists that were used in Set A, such that the lists that were first in Set A were second in Set B. The word lists that were used in each treatment combination in Set A were also rotated, as a block, to a different treatment combination in Set B. Although this materials manipulation did put different critical items into different treatment combinations, it did not provide a direct contrast on the same materials between the experimental manipulation. (This necessary control condition was implemented in Experiment 2). Finally, each of the four main experimental conditions was replicated six times within each participant, and the "critical item" data are proportions based on six observations.

In Experiment 2, the 48 lists from Experiment 1 were randomly repaired. Six such pairs were assigned to each of the four treatment combinations—Experimental Condition Test List 1, Experiment Condition Test List 2, Control Condition Test List 1, Control Condition Test List 2. In Set B the exact lists that had been used in the Experimental Conditions in Set A (except that the critical words were now replaced by random, unrelated words) were used in the corresponding Control Conditions, and the exact lists that had been in the two Control Conditions in Set A were used in the corresponding Experimental conditions in Set B. Thus, if in Set A, in some trial the "sweet" list had been presented first and the "chair" list second in the Experimental Condition Test List 1 (such that the word "sweet" was presented in the second list), the related trial in Set B would again have the "sweet" list presented first, the "chair" list presented second, and the first list would again be tested. However, now this trial would be in the Control Condition, and the critical word "sweet" would *not* be presented in the study lists.

Materials. Twenty-four lists were constructed using the materials of Roediger and McDermott (1995), and 24 additional lists were formed from the University of Florida associate norms database (Nelson, McEvoy, & Schreffier, 1994). Each list consisted of the seven highest associates of a critical item. Two unrelated words were also inserted into the list. In half of the lists, one of the unrelated words was the critical item. This item always occupied one of the middle three positions in the 7-item list, in Experiment 1, and was on average in position 6 in Experiment 2. In the computer-generated experiment List 1 words were always presented in blue at the top of the screen, and were surrounded by a thick oval border. List 2 words were in red and in a different font from the List 1 words. They appeared at the bottom of the screen and were surrounded by a rectangular border. In Experiment 2, List 1 items were always presented on blue cards and List 2 items on pink cards. These contextual differences were intended to help participants discriminate between the two lists.

In Experiment 1, the test lists consisted of 12 words: the critical item for both lists, two old List 1 words, two old List 2 words, a new word that was a low associate (the 8th-ranked associate) of the critical item related to List 1, a new word that was a low associate of the critical item related to List 2, and four new unrelated words. The order of these words within the test list was randomised across trials and across participants. In the yes/no task, List 1 was assigned to be called "old" in half the trials, and List 2 was "old" in the other half. Counterbalancing across all the within-particpants factors was maintained during this assignation.

In Experiment 2, eight words on each trial were used as test probes: the critical items from Lists 1 and 2, two presented exemplars from each of the two lists, and two unrelated new words.

Procedure. In Experiment 1 the trials were presented on a computer monitor. Participants received the following instructions verbally:

You will be doing the same task over 24 trials. For each trial, this is what you will do: First, click the mouse button to begin the trial. Then you will see a list of words appearing at the top of the screen. They will be in blue, and will appear one by one, rather quickly. After the list of blue words is over, you will see a number in the middle of the screen. Count backwards, and aloud, in steps of 3 from this number. For example, if you see 100, you would say, "97 ... 94 ... 91... and so on." Then the red list will begin. Words will appear at the bottom of the screen, one by one, in red. After this list is over, you will do the counting backwards task again. Then you will see a third list of words. Instructions on the screen will inform you about what to do. In half the trials, you will have to click on the "Yes", button if you saw the word in the blue list, and click "No" otherwise. In the other half of the trials, you will have to click on the "Yes" button if you saw the word in the red list, and click "No" otherwise. Remember, the instructions on the screen will inform you about which list you are supposed to say "Yes" to. After this test is over, you can begin the next trial.

Then they went through a practice trial before beginning the 24 experimental trials. Participants began each trial by clicking on a mouse button. List 1 then appeared, one word at a time, at a rate of 2.5sec per word. At the end of the list, participants saw a number in the middle of the screen and counted backwards aloud from this number in steps of three for 16 seconds. This task was meant to minimise rehearsal of the list they had just studied. Then List 2 started, and the words appeared at the same rate as in List 1. The counting backwards task was repeated after this list, and then the instructions for the test appeared on the screen.

The instructions for recognition were: "For each of the following words, click on the 'yes' button only if the word previously appeared in the blue [red] list. Otherwise click 'no'." Blue was always List 1 and red was always List 2. These instructions appeared in blue if they were for List 1, and in red if they were for List 2. The test words appeared one at a time and participants used the mouse to choose between "yes" and "no." Participants could then begin the next trial, or take a break for a few seconds.

The participants performing the source judgement task received these instructions: "Please indicate in which list(s) the following words appeared." They then used the mouse to click on one of four buttons as each test word was presented. The buttons were labelled "Blue only", "Red only", "Blue and Red", "Neither."

The procedure was similar in Experiment 2, except that subjects were only asked for recognition. At the end of List 2, on each trial, they were presented with the instruction card indicating which list they were to remember.

Experiment 1 Results

Recognition. The main finding of interest was that the presentation of the critical item was found to inhibit false memories to that item. The main effect of Presentation of Critical Item [$F(1,29) = 9.61$, $MSe = 0.39$], was significant. There was also a main effect of Test List, [$F(1,29) = 7.85$, $MSe = 0.46$], such that the false memory effect was stronger when List 2 was tested than when List 1 was tested. These main effects were qualified by the interaction between them [$F(1,29) = 4.52$, $MSe = 0.36$], such that participants were more likely to show false memory suppression, attributable to presentation of the critical item, when List 1 was tested than when List 2 was tested. This interaction is shown in Fig. 4.

In half of the trials, when the critical item was actually presented, the recognition instructions tested not the Inducing List, but the Other List. The person would, under these instructions, be correct to call the presented critical item old. Under these conditions, the probability of calling the critical item old, when it occurred in List 2, was 0.60. When it occurred in List 1 the probability of calling the critical item old in this condition was .40. The difference between the two conditions was significant [$t(32) = 4.61$].

This result may simply indicate that items appearing in List 2 were better remembered—enjoying a recency effect. However, there was no analogous difference between other old items presented in either List 1 or List 2, as is shown in Table 1. Thus, the simple recency effect explanation of the difference in the memorability of the critical items seems unlikely. The alternative, and more active, explanation for this difference is that after having experienced a large number of related words in one list, participants may have experienced the critical item, presented afterwards in a list that had a different theme, as an unexpected solution to a puzzle. Indeed, some participants reported to the experimenter, at time of debriefing, that they experienced surprise at the critical item word having been sometimes presented in the wrong list—after the list in which its compatriots were presented—and especially noted its occurrence.

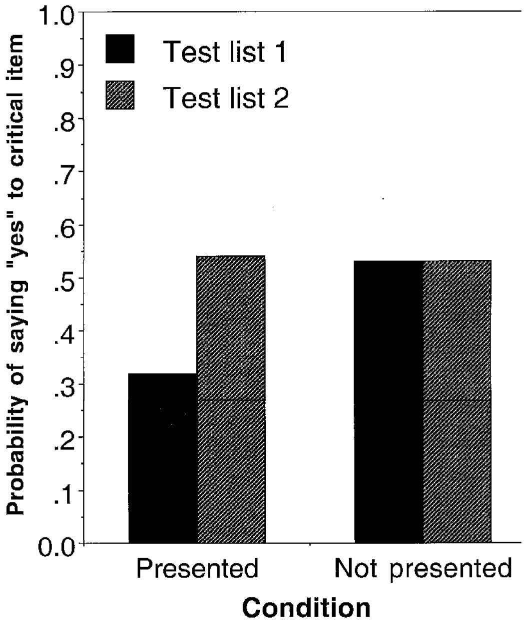

Fig. 4. *Proportion of "yes" responses to the critical item in the recognition task in Experiment 1, as a function of Test List.*

The enhanced memory for the second list critical items might be related to the false memory inhibition effect in this experiment. It is notable that the critical items that were remembered better were also the items that produced false memory inhibition. Perhaps it is necessary that the actual presentation of the critical item itself remembered crisply, (presumably also with the information that it occurred in a context different from the inducing list) if it is to be useful in inhibiting the false memory.

Source Judgements. When attribution to the inducing list was considered in isolation in the source judgement data, there was a main effect of the Presence of the Critical Item, such that when the critical item was present in the alternative list there was an inhibition effect. People were less likely to say that the critical item had been in the Inducing List when it had been presented in the Other List than when it had not been presented at all [$F(1,31) = 19.58$, $MSe = 0.02$]. People were also less likely to attribute the critical item to the Inducing List if they were tested for List 1 than if they were tested for List 2 [$F(1,31) = 40.93$, $MSe = 0.03$]. There was an interaction between the Presence of the Critical Item and Test List [$F(1,31) = 4.81$, $MSe = 0.02$], such that the false memory inhibition effect due to the presence of the critical item in the alternate list was larger when the test list was List 1 than when it was List 2. These results are shown in Fig. 5.

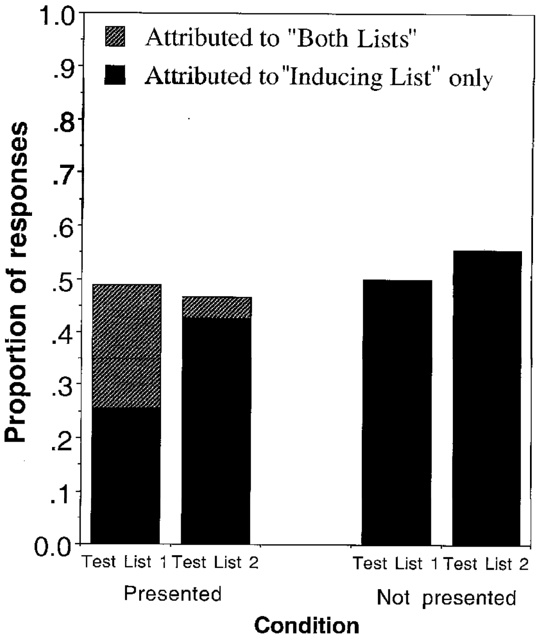

Fig. 5. *Proportion of false attribution of the critical item in the source judgement task.*

Figure 5 also shows the proportion of attributions to "both lists." People never made an attribution to both lists in the control conditions, that is, when the critical item was not presented in the alternative list. The main effect of Presence of the Critical Item was significant [$F(1,31) = 35.98$, $MSe = 0.01$]. The attribution to both lists was greater when the first list was tested than when the second list was tested [$F(1,31) = 52.92$, $MSe = 0.01$]. In addition, the interaction between the Presence of the Critical Item and the Test List was significant [$F1,31) = 37.66$, $MSe = 0.01$].

As can be seen by comparing Fig. 4 and 5, the "both" judgements in the source judgement task did not seem to translate into what seems like the logical affirmation that the critical item *had* in fact been presented in the first list, as well as the second, in the recognition task. Instead, the items that were thought to have occurred in both lists seem to have been subtracted out of the simple recognition judgements, such that the source judgments made to "Inducing List only" correspond closely to the recognition data. It is as if when participants felt, in the source monitoring task, that the critical item had been present in both lists, this assessment led to the conclusion, in the recognition task, that the item had *not* occurred in the first list.

Other Test Items. Table 1 gives the attribution to the Inducing List for the other five test items: the old words from List 1 and List 2, the extra list low associates of List 1 and List 2, and the new words. Tables 2 and 3 show the attribution of old List 1 words and old List 2 words, respectively, to each of the four possible choices in the source judgment task. New items were rejected (attributed to Neither) at levels ranging from .97 to .99 in all conditions.

Experiment 2 Results

As was the case in Experiment 1, in Experiment 2 there was a significant false memory inhibition effect such that presenting the critical items in the alternate list decreased the chance that people would false alarm to those same critical items as members of the Inducing List [$F(1,24) = 9.08$, $MSe = 0.41$]. However, this inhibition effect was dependent upon interactions with the set of items used. There was an effect of Test List [$F(1,24) = 9.51$, $MSe = 0.45$], an interaction between Set and Presentation of the Critical Item [$F(1,24) = 4.73$, $MSe = 0.21$], and a large and significant three-way interaction among Set, Test List, and Presentation of the Critical Item [$F(1,24) = 23.05$, $MSe = 0.55$]. As is shown in Fig. 6, the words in Set B exhibited an inhibition pattern like that found in Experiment 1. However, as is also shown in Fig. 6, Set A produced a different pattern, unlike that found in the first experiment.

To make the interpretation of the results clearer, we rearranged the presentation of the triple interaction data to show what happened with each of the fours groups of word lists, first when presented in the Experimental Condition (with one group of participants) then in the Control Condition (with other participants). A listing of the 24 critical words, divided into these 4 groups, are given on the abscissa of Figure 7. As can be seen from the data given in the figure, the first two word groups showed the inhibition effect when the critical word

Table 1. *Proportion of Attributions to the Inducing List in the Yes/No Recognition Task (Expt. 1)*

	Experimental Condition			
	Inducing List is List 1		*Inducing List is List 2*	
Test Item	*Critical Item Absent*	*Critical Item Present*	*Critical Item Absent*	*Critical Item Present*
Old List 1 words	.68	.74	.09	.07
Old List 2 words	.10	.09	.67	.73
List 1 low associate	.05	.06	.02	.02
List 2 low associate	.02	.02	.04	.05
New words	.04	.02	.01	.02

Table 2. *Attribution Fate of Old List 1 Words in the Source Judgement Task (in Proportions) (Expt. 1)*

	Experimental Condition			
	Inducing List is List 1		*Inducing List is List 2*	
Test Item	*Critical Item Absent*	*Critical Item Present*	*Critical Item Absent*	*Critical Item Present*
List 1	.77	.79	.79	.79
List 2	.03	.03	.02	.03
Both	.00	.00	.01	.00
Neither	.20	.17	.18	.17

Table 3. *Attributions of Old List 2 Words in the Source Judgement Task (in Proportions) (Expt. 1)*

	Experimental Condition			
	Inducing List is List 1		*Inducing List is List 2*	
Test Item	*Critical Item Absent*	*Critical Item Present*	*Critical Item Absent*	*Critical Item Present*
List 1 only	.04	.04	.06	.07
List 2 only	.77	.79	.74	.75
Both lists	.00	.01	.00	.02
Neither list	.19	.16	.21	.16

was presented, whereas the last two groups did not. We have spent many hours pondering the reasons for this particular breakdown, and have found no obvious explanation. It is notable, though, that Roediger et al. (1998) reported that there are large and apparently idiosyncratic differences among particular words in the probability of production of the basic false memory effect. At the present time, we know of no coherent explanation for either of these findings.

In summary, then, in both Experiment 1 and Experiment 2 we found inhibition of the false memory effect when the critical word was presented in an alternate list. However, the effect is conditional and rather fleeting—occurring in the first experiment only when the critical word came after the Inducing List, and in the second experiment only for a subset of the word lists. It is possible that the appearance of the inhibition effect depends on the salience of the presented critical word, as well as

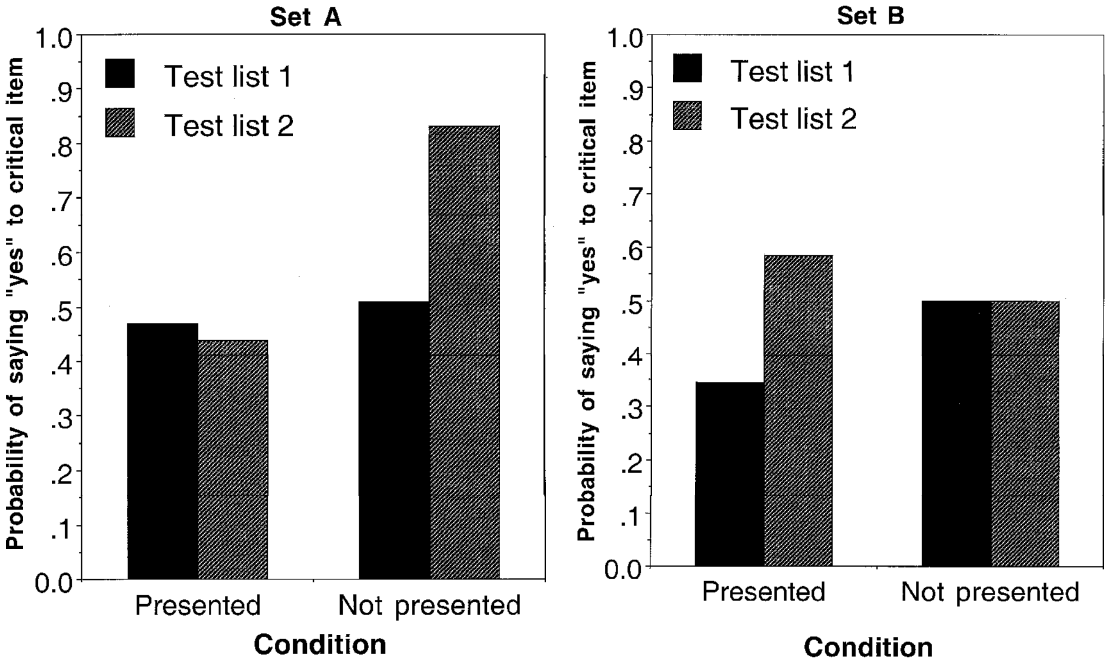

Fig. 6. *Proportion of "yes" responses to the critical item in the recognition task in Experiment 2, as a function of Set and Test List.*

on the participant registering that it is out of the context of the Inducing List. This conjecture, while appealing, would require further empirical investigation.

CONCLUSION

People appear, at least under some conditions, to be able to use source information to sharpen and improve their memories in both of the misremembering situations investigated in this article. In the misleading information paradigm, when asked, they are able to use source information to indicate that the misleading information did not actually occur in the event in question. Similarly, in the false memory paradigm, in which a memory is induced by a number of related items, the salient presence of the real critical item in a context other than that of the inducing list itself can sometimes be used to suppress false alarms to that critical item that is only induced (see Johnson & Raye, 1981; and Johnson, Kounios, & Reeder, 1994 for discussions of reality monitoring). In the latter case, the exact circumstances both of when the critical item is presented, and of which particular words are used, appear to be important. Even so, it seems that the highly salient and memorable occurrence of the critical item in the wrong context may allow people to disconfirm it as a member of the inducing list.

A phenomenon similar to the one observed here—that people reject an item based on the fact that they remember it as having occurred in the wrong context—can be found in the Jacoby (1991) exclusion condition of the process dissociation paradigm. In that paradigm, participants were explicitly told that if they remember an item as having occurred in a particular list it means that it definitely did not occur in the other list. Although the participants in our experiment were not told this, and did, in fact, subscribe incorrectly to the double presentation of the critical item in two lists when they were asked for source judgements, the paradigm we used did conform, in some respects, to Jacoby's process dissociation paradigm (Jacoby, 1991; Jacoby & Kelley, 1992; Jacoby, Toth, &

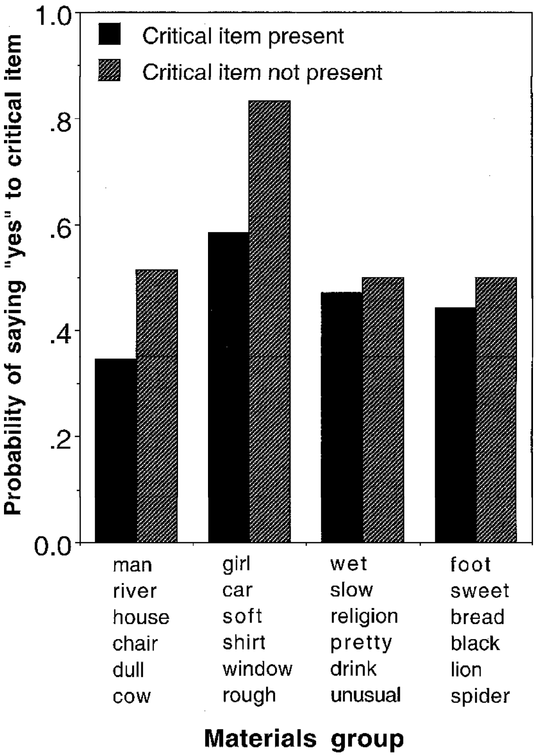

Fig. 7. *The inhibition effect in Experiment 2, as a function of the particular lists used.*

Yonelinas, 1993; Yonelinas & Jacoby, 1996), insofar as items never *actually* occurred in both lists.

It is possible that participants, under our recognition instructions, used the exclusion logic to their advantage in this experiment—deciding that an item was not in the inducing list if he or she remembered that it was in the alternate list. The source monitoring data seem to weigh against this explanation, at least on the face of it. People, in the source judgement task, freely told us that they thought (incorrectly) that many of the critical items had been presented in *both* lists. Had they been using a strict exclusion rule, even on the source judgement task, they presumably would not have subscribed to the "both" option. On the other hand, participants often mentioned at time of debriefing that, although it had seemed as if certain items that had been presented in both lists, they had responded under the hypothesis that this phenomenology was illusory, and hence they had used an exclusion rule. We could easily model the exclusion rule in CHARM, and hence account for the data in these experiments. Such a rule would have something of a post hoc quality, given the instructions in the experiments, but even so, if it is the rule people were using, it would be entirely reasonable to include it in the model in this situation.

REFERENCES

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

Belli, R.F. (1989). Influences of misleading postevent information: Misinformation interference and acceptance. *Journal of Experimental Psychology: General, 118,* 72–85.

Ceci, S.J., Ross, D.F., & Toglia, M.P. (1987). Suggestibility of children's memory: Psycholegal implications. *Journal of Experimental Psychology: General, 117,* 201–203.

Chandler, C.C. (1989). Specific retroactive interference in modified recognition tests: Evidence for an unknown cause of interference. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 256–265.

Christiaansen, R.E., & Ochalek, K. (1983). Editing misleading information from memory: Evidence for the coexistence of original and postevent information. *Memory & Cognition, 11,* 467–475.

Deese, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *Journal of Experimental Psychology, 58*(1), 17–22.

Galton, F. (1878). Composite portraiture. *Nature, 18,* 97–100.

Hintzman, D.L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review, 95,* 528–551.

Israel, L., & Schacter, D.L. (1997). Pictorial encoding reduces false recognition of semantic associates. *Psychonomic Bulletin and Review, 4,* 577–581.

Jacoby, L.L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30,* 513–541.

Jacoby, L.L., & Kelley, C.M (1992). A process-dissociation framework for investigating unconscious influences: Freudian slips, projective tests, subliminal perception, and signal detection theory. *Current Directions in Psychological Science, 1,* 174–179

Jacoby, L.L., Toth, J.P., & Yonelinas, A.P. (1993). Separating conscious and unconscious influences of memory: Measuring recollection. *Journal of Experimental Psychology: General, 122,* 139–154.

Johnson, M.K., Hashtroudi, S., & Lindsay, D.S. (1993). Source monitoring. *Psychological Review, 114,* 3–28.

Johnson, M.K., Kounios, L, & Reeder, J. (1994). Time-course studies of reality monitoring and recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 1409–1419.

Johnson, M.K., & Raye, C.L. (1981). Reality monitoring. *Psychological Review, 88,* 67–85.

Kensinger, E.A., & Schacter, D.L. (this issue). When true memories suppress false memories: Effects of ageing. *Cognitive Neuropsychology, 16,* 399–415.

Koutstaal, W., & Schacter, D.L. (1997). Gist-based false recognition of pictures in older and younger adults. *Journal of Memory and Language, 37,* 555–583.

Kroll, N.E., & Timourian, D.A. (1986). Misleading questions and the retrieval of the irretrievable. *Bulletin of the Psychonomic Society, 24,* 165–168.

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. (1989). The eyewitness suggestibility effect and memory for source. *Memory & Cognition, 17,* 349–358.

Loftus, E.F. (1979). *Eyewitness testimony*. Cambridge, MA: Harvard University Press.

Loftus, E.F., & Hoffman, H.G. (1989). Misinformation and memory: The creation of new memories. *Journal of Experimental Psychology: General, 118,* 100–104.

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

McDermott, K.B. (1996). The persistence of false memories in list recall. *Journal of Memory and Language, 35,* 212–230.

Metcalfe, J. (1990). A composite holographic associative recall model (CHARM) and blended memories in eyewitness testimony. *Journal of Experimental Psychology: General, 119,* 145–160.

Metcalfe, J. (1991a). Recognition failure and the composite memory trace in CHARM. *Psychological Review, 98,* 529–553.

Metcalfe, J. (1991b) Composite memories. In W. Hockley & S. Lewandowsky (Eds.) *Relating theory and data. Essays on human memory in honor of Bennet B. Murdock* (pp. 399–423). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Metcalfe, J. (1993a). Novelty monitoring, metacognition, and control in a composite holographic associative recall model: Implications for Korsakoff amnesia. *Psychological Review, 100,* 3–22.

Metcalfe, J. (1993b). Monitoring and gain control in an episodic memory model: Relation to P300 event-related potentials. In A.F. Collins, S.E. Gathercole, M.A. Conway, & P.E. Morris (Eds.), *Theories of Memory* (pp. 327–354) Hove, UK: Lawrence Erlbaum Associates Ltd.

Metcalfe, J. (1997). Predicting syndromes of amnesia from a Composite Holographic Associative Recall/Recognition Model (CHARM). *Memory, 5,* 233–253

Metcalfe, J., Cottrell, G.W., & Mencl, W.E. (1992). Cognitive binding: A computational-modeling analysis of a distinction between implicit and explicit memory. *Journal of Cognitive Neuroscience, 4,* 289–298.

Metcalfe Eich, J. (1982). A composite holographic associative recall model. *Psychological Review, 89,* 627–661.

Metcalfe Eich, J. (1985). Levels of processing, encoding specificity, elaboration, and CHARM. *Psychological Review, 91,* 1–38.

Nelson, D.L., McEvoy, C.L., & Schreiber, T.A. (1994). *The University of South Florida word association, rhyme and word fragment norms.* Unpublished manuscript.

Norman, K.A., & Schacter, D.L. (1997). False recognition in younger and older adults: Exploring the characteristics of illusory memories. *Memory & Cognition, 25,* 838–848.

Payne, D.G., Toglia, M.P., & Anastasi, S. (1994). Recognition performance-level and the magnitude of the misinformation effect in eyewitness memory. *Psychonomic Bulletin and Review. 1,* 376–382.

Reyna, V.F., & Brainerd, C.J. (1995). Fuzzy-trace theory: An interim synthesis. *Learning and Individual Differences, 7,* 1–75.

Robinson, K.J., & Roediger, H.L. (1997). Associative processes in false recall and false recognition. *Psychological Science, 8,* 231–237.

Roediger, H.L., Gallo, O.A., Watson, J.M., & Balota, D.A. (1998). *Eliminating and enhancing illusory memories.* Presentation given at the Annual Meeting of the Psychonomic Society, Nov. 21, Dallas, Texas.

Roediger, H.L., & McDermott, K.B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory. and Cognition, 21,* 803–814.

Schacter, D.L., Israel, L., & Racine, C. (1999). Suppressing false recognition in younger and older adults: The distinctiveness heuristic. *Journal of Memory and Language, 40,* 1–24.

Schacter, D.L., Kaszniak, A.W, Kihlstrom, J.F, & Valdiserri, M. (1991). The relation between source memory and aging. *Psychology and Aging, 6,* 559–568

Schacter, D.L., Verfaellie, M., Anes, M. D., & Racine, C. (1998). When true recognition suppresses false recognition: Evidence from amnesic patients. *Journal of Cognitive Neuroscience, 10,* 668–679.

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.

Yonelinas, A.P., & Jacoby, L.L. (1996). Response bias and the process-dissociation procedure. *Journal of Experimental Psychology: General, 125,* 422–434.

Zaragoza, M.S., & Koshmider, J.W. (1989). Misled subjects may know more than their performance implies. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 246–255.