CONSCIOUSNESS AND COGNITION **5,** 176–196 (1996)
ARTICLE NO. 0011

Continuing Influences of To-Be-Forgotten Information

ELIZABETH LIGON BJORK[1] AND ROBERT A. BJORK

Department of Psychology, University of California, Los Angeles, California 90095-1563

In the present paper, we first argue that it is critical for humans to forget; that is, to have some means of preventing out-of-date information from interfering with the recall of current information. We then argue that the primary means of accomplishing such adaptive updating of human memory is retrieval inhibition: Information that is rendered out of date by new learning becomes less retrievable, but remains at essentially full strength in memory as indexed by other measures, such as recognition and word-fragment completion. We conclude with a speculation that certain unconscious influences of prior events may, in fact, be stronger if those events were to be forgotten rather than to be remembered.

Many situations in the real world cue us—directly or indirectly—to forget. At one time or another, we have all probably said, or had said to us, expressions like the following: ''Oh, forget that. I read the wrong address. Here's the correct one.'' Or, ''Forget that! Here's a better way to get there.''

Such explicit situations are greatly outnumbered by situations where a cue to forget is implicit, but nonetheless clear. Everyday illustrations are the need to remember our current phone number—not ones we have had in the past; where we parked the car this morning—not yesterday morning; what the trump suit is on this hand—not the last one; how to select and move text in this program—not the one on our old computer; and so forth. In the laboratory, a number of our research paradigms, such as the Brown-Peterson paradigm and paired-associate list-learning procedures of the A–B, A–C variety, involve an implied—but clear—cue to forget. When each successive word trigram (or some other set of to-be-remembered items) is presented in the Brown-Peterson procedure, for example, it is clear to subjects that the preceding trigram should now be forgotten—that continuing to remember that trigram, now out of date, is a source of confusion and intrusions. Similarly, when—at the start of A–C learning—it is clear that the List-1 B responses are now history and a potential source of errors, there is again a clear, if implicit, cue to forget.

Thus, in a great variety of situations, on a great range of time scales, we need to update our memories. We need some means to set aside, segregate, erase, or inhibit out-of-date information and replace it with current updated information. The present paper addresses the issues of how such updating is accomplished and the fate of the information we have intentionally forgotten in order to do so—particularly the question of its continued existence or strength in memory and its continuing influences—or lack thereof—on our ongoing thoughts and actions, even though we may be unable, consciously, to recall that information. Although a cue to forget can often elimi-

[1] To whom reprint requests should be addressed at Department of Psychology, University of California, Box 951563, Los Angeles, CA 90095-1563.

1053-8100/96 $18.00

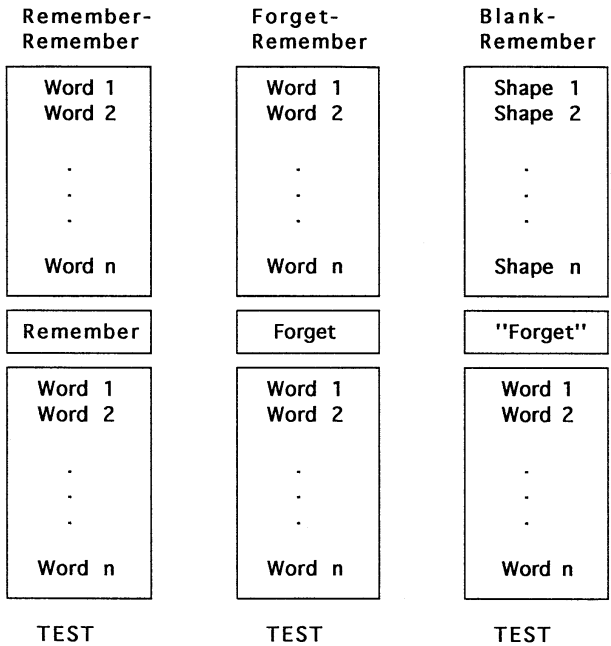

FIG. 1. Illustration of the typical paradigm used in the list method of directed forgetting.

nate entirely the proactive interference that would normally be attributable to the to-be-forgotten information, as demonstrated in the studies to be presented below, there is clear evidence using other indices that the to-be-forgotten items remain in memory. There is also evidence, however, from studies of both hypnotic directed forgetting (see, for example, Kihlstrom, 1983) and non-hypnotic directed forgetting (see Geiselman, Bjork, & Fishman, 1983) that the state of such to-be-forgotten information in memory has been altered. In terms of source attribution and measures of subjective organization, for example, the state of to-be-forgotten items in one's memory appears to differ markedly from the state of corresponding to-be-remembered items.

BACKGROUND FOR THE PRESENT RESEARCH

Basic Procedures

In laboratory settings, when subjects are cued to forget or to update their memories, they can, in fact, do so. Figure 1 illustrates a typical paradigm used to study memory updating or directed forgetting in the laboratory, sometimes referred to as the list method of directed forgetting. In the typical list method of directed forgetting, subjects are presented with lists of items to remember, with the items (e.g., common words) presented one at a time. After some number of items (typically, half) have been presented, the list presentation is interrupted with a cue either (a) to forget the preceding items, as illustrated in the middle list of Fig. 1, or (b) to keep on remembering the preceding items, as illustrated in the left list of Fig. 1. In the following sections of the present paper, these two types of list or cueing conditions are designated as the forget–remember list type or condition and the remember–remember list type or condition, respectively. It is important to realize in this paradigm that during the presentation of the first part of a list that is to be a forget–remember list, subjects have no way of anticipating that a mid-list cue to forget will be presented. Thus, for both types of lists, precue items must be processed by subjects in the same way—that is, in both cases, subjects must be trying to learn the items for a possible later memory test.

In this general type of directed-forgetting paradigm, a control list or condition, as illustrated by the right list in Fig. 1, is sometimes employed. In this type of list, referred to as a blank–remember list or control condition, no words are presented to subjects in the equivalent first half of the list. Instead, subjects see and process other types of information. For example, subjects may be asked to perform some sort of discrimination task on successive pairs of shapes during this part of the list, as indicated in Fig. 1.

One final aspect of the present basic paradigm needs to be clarified, as it seems to be a source of some confusion. Namely, it needs to be made clear that this basic paradigm can be utilized in a totally within-subjects design, as, indeed, was done in all of the studies reported in the present paper. That is, all three of the list types described above can be repeatedly presented to the same subject with no diminution in the effectiveness of the forget cue. One might well think that a cue to forget can only be introduced once—as a surprise, so to speak. For example, some studies using the list method of directed forgetting introduce the mid-list cue to forget as a surprise to subjects, either by indicating after the presentation of the first list that the first list was just for practice and can be forgotten or, in what is sometimes referred to as the ''whoops!'' procedure, by indicating after presentation of the first list that it was somehow incorrectly presented to them, that they should thus forget it and learn the correct list, which will be presented next. Then, following presentation of the second list, subjects are tested on the items they were instructed to forget as well as those they were instructed to remember. Obviously, such misleading use of the forget cue can only be administered once for any given subject, and, thus, forget–remember versus remember–remember list conditions must be manipulated as a between-subjects variable in such directed-forgetting studies.

However, it is also the case that subjects can be instructed prior to the presentation of any lists that if a forget cue occurs, they will not be tested for their memory of the preceding items, only for their memory of the items to follow the forget cue; whereas, if a remember cue occurs, they will be tested on their memory for the preceding items as well as for the items to follow the remember cue. In all of the studies reported in the present paper, subjects were fully informed as to the meaning of the forget and remember cues before any lists were presented, and the validity of the forget instruction was never violated. Consequently, all of the different types of lists could be repeatedly presented to the same subjects. It should be emphasized, however, that—although informed that cues to forget and cues to remember can occur on any given list—subjects have no way of predicting when a cue to forget or when a cue to remember will be presented and, thus, up to the point of the cue, subjects must process all presented information in the same way—that is, as though they will be asked to remember it. (See also Bjork, 1970, 1978, and Reitman, Malin, Bjork, &

Higman, 1973) for other examples of paradigms using repeated presentations of the cue to forget in within-subjects designs.)

Finally, we need to mention that there is another frequently used paradigm for studying directed-forgetting effects, sometimes referred to as the item-by-item cueing procedure, in which the presentation of each individual item is followed by a cue either to forget or to remember that item. Although there are a number of similarities between the effects obtained with these two methods, there are also compelling reasons to believe that the underlying processes involved in producing these effects may be quite different. Thus, the treatment of directed-forgetting effects in the present paper is confined to a consideration of those effects obtained in the list method of directed forgetting.

Recall Phenomena

Extrapolating from many previous studies using the above or similar types of list-method directed-forgetting paradigms, directed-forgetting results can be summarized in terms of three basic and robust findings. First, the recall of postcue items (i.e., items presented after the mid-list cue) is better for forget–remember lists than for remember–remember lists. That is, the recall of postcue items in the forget–remember condition suffers less evidence of proactive interference from precue words than does the recall of postcue items in the remember–remember condition.

Second, compared to the blank–remember or control condition where no items are presented in the first half of the list, recall of postcue items in the forget–remember condition often does not differ from recall of postcue items in the blank–remember condition. That is, recall of postcue items shows no evidence of suffering from any proactive interference owing to the to-be-forgotten items that were presented in the first half of the list. The subjects' recall performance is as though the to-be-forgotten items had not been presented.

Third, if subjects' ability to recall items that they were instructed to forget is unexpectedly tested (obviously, such a violation of prior instructions can only occur once for any given subject), the recall of such to-be-forgotten items is severely depressed compared to equivalent items that subjects were instructed to remember.

To summarize, three basic and robust findings obtained using the list method of directed forgetting are that (a) recall of postcue items in the remember–remember type of list is poorer than recall of postcue items in the forget–remember type of list, which (b) often does not differ from recall of postcue items in the control or blank–remember type of list, where no precue items to be learned were presented at all. And (c), recall of precue items in the forget–remember type of list—if tested—is severely depressed compared to recall of precue items that were to be remembered. (See Bjork, 1970, 1972, 1989; and Johnson, 1994, for a discussion and review of studies reporting these basic directed-forgetting phenomena.)

Retrieval Inhibition as the Primary Mechanism in Directed Forgetting

For a variety of reasons, the explanation of these directed-forgetting effects that we have come to favor is in terms of retrieval inhibition. That is, when subjects are told to forget preceding information and then presented with new information to learn, a process is initiated that *inhibits* the subsequent *retrieval* of the to-be-forgotten information. Because such to-be-forgotten items are not retrievable, they do not interfere with recall of current to-be-remembered information. Furthermore, while this updating process inhibits the retrieval of the to-be-forgotten information, it leaves its strength in memory—as indicated by other measures—unaffected. (See also Bjork, 1989, for a further review and discussion of the role of retrieval inhibition in the production of directed forgetting effects.)

Evidence for this last assumption—that the strength in memory of to-be-forgotten information is not affected by the instruction to forget it—comes from the following findings: (a) Recognition performance for to-be-forgotten items is unimpaired (e.g., Block, 1971; Elmes, Adams, & Roediger, 1970; Geiselman et al., 1983; Gross, Barresi, & Smith, 1970); (b) in a relearning paradigm, to-be-forgotten items are relearned as readily as to-be-remembered items (e.g., Reed, 1970; Geiselman & Bagheri, 1985); and (c) the proactive interference of precue items that is eliminated by the instruction to forget can—under certain circumstances—be reinstated at full strength. The first demonstration of this last effect was reported by Bjork, Bjork, and Glenberg (1973). We review their study in some detail below as the procedures they employed are quite similar to those used in the present Experiments 1 and 2, and because it was the implications of this finding for the type of retrieval inhibition involved in the production of directed-forgetting results that largely motivated the present research.

In the Bjork et al. (1973) study, which employed the list method of directed forgetting, all subjects were presented with each of the types of list illustrated in Fig. 1: remember–remember, forget–remember, and blank–remember. Subjects' ability to recall to-be-remembered items, however, was tested either by an immediate free recall test or by a free recall test delayed by one of two types of intervening tasks: an arithmetic task or a recognition task, in which a subset of to-be-forgotten items appeared as distractors. For either the remember–remember or forget–remember lists, 16 randomly selected common words, presented for 3 s each by means of a slide projector, preceded and followed the mid-list cue to forget or to remember, which was a row of minus signs or a row of plus signs, respectively, and also presented for 3 s. For the blank–remember lists, the precue items were 16 pairs of simple shapes, also presented at 3 s per pair, for which subjects had to report—on a response sheet—the number (0 to 3) of dimensions (shape, size, color) on which the members of each pair differed.

Each subject received each type of list three times, with each type being tested once in each of the three testing conditions: an immediate recall test of all to-be-remembered words that had been presented; a recall test of all to-be-remembered words delayed by an arithmetic task; and a recall test of all to-be-remembered words delayed by a forced-choice recognition test. These different types of tests were administered by way of a booklet that was handed out before the start of each list, and the subject was allowed to proceed through the test booklet at his or her own pace, but was not allowed to return to any previous page or task. Additionally, booklets were constructed in such a way that subjects could not tell which type of test would be given until instructed to open the booklet after the last word of the current list had been presented.

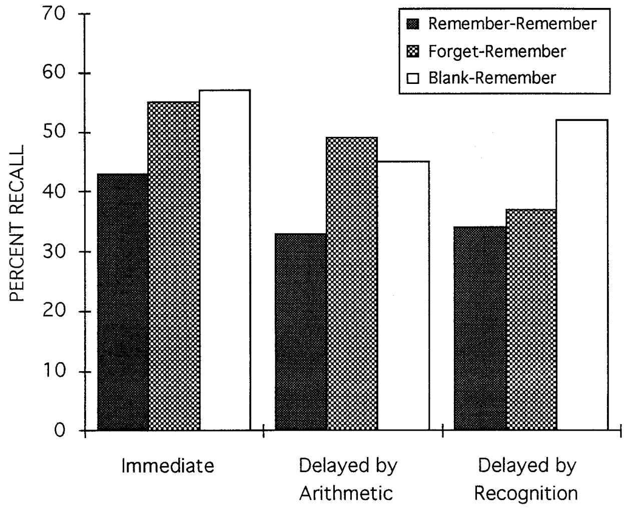

FIG. 2. Percentages of correct recall for postcue to-be-remembered words that had *not* appeared on the recognition tests for each list type and testing condition.

In the recognition test, eight pairs of words were presented and subjects were asked to circle the word in each pair that had been presented in the second half of the list—that is, after the mid-list cue. Thus, for all three list types, the correct recognition response was always a word that subjects had been instructed to remember. For remember–remember lists, four of the pairs on the recognition test contained new words as the distractor items (i.e., words that had not been presented previously in the experiment) and four of the pairs contained words from the first half of the list as the distractor items. For forget–remember lists, four of the pairs on the recognition test contained new words as the distractors and four pairs contained words from the first half of the list as distractors—that is, to-be-forgotten words. For blank–remember lists, all distractor items were new words. The eight ''old'' words taken from the second half of the list to appear on the recognition test as the correct items were randomly selected with the constraint that only two could be taken from each quarter of the list; similarly, the four ''old'' words taken from the first half of the list to appear as foils on the recognition test were randomly selected with the constraint that only one word could be taken from each fourth of the list. The position of pairs containing old versus new distractor items were spread more or less evenly across the recognition test.

Finally, across subjects, lists were counterbalanced such that each specific list of words occurred equally often as each type of list, each specific list half occurred equally often as the first and second half of a total list, and each combination of list type and testing condition appeared equally often in all presentation positions. Thus, there are no item differences or presentation-order differences involved when making comparisons across list-type conditions, testing conditions, or between forget and remember words.

The results obtained by Bjork et al. (1973) are presented in Fig. 2 where the correct recall percentages observed for postcue to-be-remembered words that had *not* appeared on the recognition tests are shown for each type of list in the three types of recall conditions (immediate test; test delayed by arithmetic task; test delayed by recognition test). An examination of the three left bars, which display the results for the immediate recall test, indicates that the basic directed-forgetting findings were obtained. Indeed, planned comparison tests revealed the percentage of correct recall for postcue to-be-remembered words to be significantly poorer in the remember–remember list condition than in the forget–remember list condition; whereas, recall of postcue to-be-remembered words in the forget–remember list condition did not differ significantly from that in the control or blank–remember list condition, where there were no words presented prior to the mid-list cue.

A comparison of the immediate recall results to the results obtained when recall was delayed by an arithmetic task—the middle three bars—reveals, first, that forgetting did occur over the retention interval as indicated by the lower overall levels of recall performance, but, more importantly, that the pattern of recall across the three list-type conditions did not appear to change. That is, there is no evidence of spontaneous recovery of proactive interference owing to the to-be-forgotten items in the forget–remember lists during the retention interval filled with an arithmetic task. A planned-comparison analysis of these results confirmed this picture, revealing that correct recall of postcue to-be-remembered words in the remember–remember list condition was again significantly poorer than correct recall in the forget–remember list condition, which did not significantly differ from correct recall performance in the control or blank–remember list condition.

An examination of the three right bars, however, indicates that the pattern of recall performance dramatically changed when recall of the postcue to-be-remembered words was delayed—not by the arithmetic task—but by the recognition test. In this situation, correct recall performance in the forget–remember list condition appears to fall to the level of that in the remember–remember list condition, with both being poorer than performance in the control or blank–remember condition. Again, a planned-comparison analysis, based only on postcue items that had *not* appeared on a recognition test, confirmed this apparent pattern, revealing the correct recall of postcue to-be-remembered words in the remember–remember and forget–remember lists not to be significantly different from one another, while both were significantly poorer than correct recall performance for the control or blank–remember lists.

A compelling explanation of the difference in the pattern of recall performance obtained between the two delayed test conditions in this study is that spontaneous recovery of the proactive interference owing to the to-be-forgotten items took place during the retention interval filled with a recognition test involving some to-be-forgotten items as distractors, but did not occur during an interval in which subjects were engaged in performing an arithmetic task. With respect to this explanation, it is interesting to note that the recognition test itself showed evidence of proactive interference owing to the to-be-forgotten items: Subjects' false alarm rate for foils drawn from the to-be-forgotten list was 19%, but only 4% for new foils.

Unlike the intervening arithmetic task, the intervening recognition test released the inhibition that had been imposed upon the precue words by the instruction to forget them, as indicated by the reinstatement of the proactive interference effects of the precue items in the latter but not the former condition. An important question arises from this dramatically different outcome of the two types of intervening tasks: Namely, what critical aspects of the recognition test lead to this release of inhibition? The recognition test necessarily involved certain types of search processes, for example. Perhaps there is something critical about the type of search and comparison processes that were evoked by the forced-choice nature of the recognition task? However, in addition, the recognition test involved the reexposure of four to-be-forgotten items as distractors. Which of these factors, if not both, is critical: the processes invoked by the recognition task, whatever they might be, or the reexposure of to-be-forgotten items? We turn now to the report of our first experiment, which attempted to answer these and other questions raised by the Bjork et al. findings.

EXPERIMENT 1: NECESSARY CONDITIONS FOR THE REINSTATEMENT OF INTERFERENCE[2]

To investigate further the issue of what factors are critical in releasing the inhibition of to-be-forgotten items, subjects were again presented with the three list types as used in the Bjork et al. (1973) study. In the present study, however, subjects' ability to recall to-be-remembered words was measured either by an immediate recall test or by a recall test delayed by two types of yes/no recognition tests: one that presented precue items as some of the foils and one that did not—that is, all foils were new items. The change to a yes/no recognition test would allow us to answer the question of whether there was something critical about the forced-choice nature of the recognition test used in the Bjork et al. study that led to the release of inhibition of to-be-forgotten items. For example, perhaps the necessity of having to compare and choose between the remember word and the forget word within a forced-choice pair initiated a type of memory search that helped to lead subjects back to the learning episode involving the to-be-forgotten items. If so, then on a yes/no recognition test in which subjects need consider only one item at a time, the inhibited state of the to-be-forgotten items might not be released. The inclusion of both a yes/no recognition task that presented precue items as foils and one that did not would allow us to answer the question of whether it is processes initiated by the recognition task or the reexposure of to-be-forgotten items that is the critical or essential factor in producing the release from inhibition.

Method

Subjects. The subjects were 18 undergraduates attending the University of California, Los Angeles, and their participation partially fulfilled an introductory psychology course requirement.

Design and procedure. All subjects were presented with the three list types illustrated in Fig. 1: remember–remember, forget–remember, and blank–remember. For either the remember–remember or forget–remember lists, 16 randomly selected common words, presented at 3 s each by means of a slide projector, preceded and followed

[2] Portions of Experiment 1 were reported in a paper by Bjork, Bjork, and White (1984, November) presented at the meeting of the Psychonomic Society (San Antonio, TX).

the mid-list cue to forget or to remember, which was a row of minus signs and plus signs, respectively, also presented for 3 s. For the blank–remember lists, the precue items were 16 pairs of simple shapes, also presented at 3 s per pair, for which subjects performed the same judgment task as in the Bjork et al. study.

Each subject received each type of list three times, with each type being tested equally often in each type of testing condition: an immediate recall test of all to-be-remembered words that had been presented; a recall test delayed by a yes/no recognition test in which precue to-be-forgotten items appeared as some of the foils; and a recall test delayed by a yes/no recognition test in which no to-be-forgotten items appeared as foils. The subjects' task on either type of recognition test was to circle any words that they recognized as having been presented after the mid-list cue (or in the second half of the list)—thus, a remember word for all list types. The different types of tests were administered by way of a booklet that was handed out before the start of each list; subjects were allowed to proceed through the test booklet at their own pace, but not allowed to return to any previous page or task, and subjects could not tell which type of test would be given until instructed to open the booklet after the last word of the current list had been presented.

Each type of recognition test contained 16 words, 8 of which were postcue items—that is, the correct ''yes'' items. On the recognition tests involving precue items as foils, 4 of the foils were precue items—that is, 4 words from the first half of the list. Thus, for the remember–remember lists, 4 precue to-be-remembered words appeared as foils, and for the forget–remember lists, 4 to-be-forgotten words appeared as foils. The remaining 4 foils, for either type of list, were new words. On the recognition test with no precue items as foils, all 8 foils were new items for both remember–remember and forget–remember lists. The recognition tests following control or blank–remember lists contained only new items as foils because no words were presented prior to the mid-list cue. Old words to appear on the recognition tests either as correct items (i.e., words from the second half of the lists) or as foils (i.e., words from the first half of the lists) were randomly selected with the following constraints: Of the 8 postcue words to appear as correct items, only 2 could be taken from each quarter of the list; of the 4 precue words to appear as foils, only 1 could be taken from each fourth of the list. The positions of correct items and foils were spread more or less evenly across the recognition test.

Finally, across subjects, lists were counterbalanced such that each specific list of words occurred equally often as each type of list, each list half occurred equally often as the first and second half of a total list, and each combination of list type and testing condition appeared equally often in all presentation positions. Thus, there are no item differences or presentation-order differences involved when making comparisons across list type conditions, testing conditions, or between forget and remember words.

Results

The correct recall percentages for only those postcue to-be-remembered words that had *not* appeared as correct ''yes'' responses on a recognition test are shown in Fig. 3 for the different types of lists and testing conditions. Looking first at the results for the immediate recall test, displayed in the three left bars of Fig. 3, recall perfor-

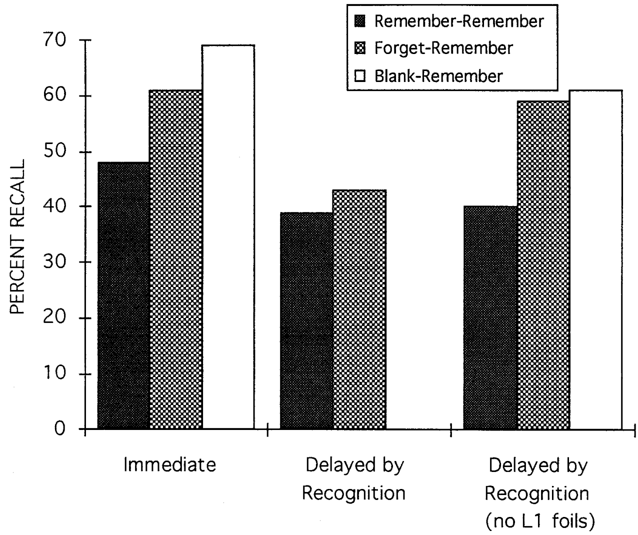

FIG. 3. Percentages of correct recall for postcue to-be-remembered words that had *not* appeared on the recognition tests for each list type and testing condition.

mance seems to replicate the pattern obtained in the Bjork et al. (1973) study. Indeed, planned comparisons performed on these data revealed the percentage of correct recall for postcue to-be-remembered words to be significantly poorer in the remember–remember list condition than in the forget–remember list condition; whereas, recall performance in the forget–remember list condition did not differ from that in the control or blank–remember list condition, where no words had been presented prior to the mid-list cue, $F(1, 17) = 6.00$, $MS_e = .03$, $p < .025$; $F(1, 17) = 2.51$, $MS_e = .02$, $p > .13$, respectively. Interestingly, not only was the pattern of immediate recall performance obtained in the Bjork et al. study replicated by the present immediate recall performance, but the levels of recall performance obtained in the two studies—conducted at different times and places, with different subjects, and using different word lists—are also remarkably similar.

Looking next at the performance obtained when recall was delayed by a yes/no recognition test that included some precue items as foils, we see a pattern that looks similar to that previously obtained by Bjork et al. when recall was delayed by a forced-choice recognition task in which some of the distractors were precue items. Namely, recall performance in the forget–remember lists appears to have decreased to the level of that for remember–remember lists, with both being less than that for the control or blank–remember lists. Planned comparisons confirmed this apparent pattern with delayed recall performance for the remember–remember lists and the forget–remember lists not being significantly different, $F(1, 17) = 0.42$, $MS_e = .04$, $p > .50$; whereas, recall performance for forget–remember and for remember–remember lists were both significantly poorer than that for the control or blank–remember lists, $F(1, 17) = 7.26$, $MS_e = .04$, $p < .015$, $F(1, 17) = 19.79$, $MS_e = .02$, $p < .0004$, respectively.

Looking finally at the performance obtained for the three list types when the recall test was delayed by a recognition test that did *not* re-present precue items as distractors—displayed by the right three bars—a dramatically different pattern emerges. The pattern here appears to reflect, once again, the basic directed-forgetting results with recall of postcue to-be-remembered words being poorer in the remember–remember list condition than in the forget–remember list condition, with the latter recall performance not differing from that of the control or blank–remember list condition. Planned comparisons also confirmed this pattern of results, revealing postcue recall for the remember–remember lists to be significantly poorer than that for forget–remember lists, whereas performance levels for the forget–remember and the blank–remember conditions did not differ, $F(1, 17) = 15.08$, $MS_e = .02$, $p < .001$; $F(1, 17) = 0.27$, $MS_e = .03$, $p > .50$, respectively.

Discussion

Comparing the results obtained across the present experiment and that of Bjork et al. (1973), then, the yes/no recognition task with precue to-be-forgotten words as foils of the present study has acted like the forced-choice recognition task of the earlier study, whereas the yes/no recognition task with no precue to-be-forgotten words as foils has acted like the arithmetic task of the earlier study. In the first case, the proactive interference owing to the to-be-forgotten items was reinstated following the recognition task, whereas in the second case, there is no reinstatement of proactive interference. A reexposure of to-be-forgotten items is, thus, an apparently necessary condition for reinstating their potential interference or for releasing their inhibition.

To summarize the findings to this point, from the combined results of these two experiments, several important extensions to the basic directed-forgetting pattern of results have emerged. First, precue to-be-forgotten items—unlike precue to-be-remembered items—are somehow prevented from interfering with the recall of postcue to-be-remembered items on either an immediate recall test or on a recall test that is delayed by an arithmetic task. When, however, the recall test for postcue to-be-remembered items is delayed by a recognition test of postcue to-be-remembered items on which a few to-be-forgotten items appear as distractors, proactive interference from the entire list of precue to-be-forgotten items is reinstated, as evidenced by the drop in recall of postcue to-be-remembered items to the level of that obtained when subjects do not receive instructions to forget the precue items. If, however, the recognition test for postcue to-be-remembered items does not re-present any precue to-be-forgotten items as distractors, then the entire set of to-be-forgotten items remains inhibited, as evidenced by the lack of any proactive interference effects.

Thus, there is no spontaneous recovery of the interference owing to the to-be-forgotten items when recall is delayed either by some dissimilar task, such as solving arithmetic problems, or by a recognition test that does not present any to-be-forgotten items as distractors. On the other hand, when to-be-forgotten items appear as distractors in an intervening recognition test for some of the postcue items, the later recall of postcue items suffers from proactive interference from the entire set of to-be-forgotten items. That is, in this latter situation, the inhibition of the entire set of to-be-forgotten items is apparently released, even though only a subset of to-be-forgotten items appear as foils on the recognition test.

Taken together, the foregoing findings raise interesting questions regarding the nature of the inhibition that is created in directed-forgetting paradigms. What exactly is inhibited as a consequence of the directed-forgetting instructions and then released by certain tasks—such as the recognition test involving to-be-forgotten items as distractors? One possibility is that the overall access to a to-be-forgotten item as a lexical entry is inhibited—or possibly, something more purely episodic is being inhibited, such as the process of retrieving items from the specific to-be-forgotten episode. These various possibilities as to the nature of the inhibition created by instructions to forget are addressed in the next experiment of the present paper by examining the effects of to-be-forgotten information when assessed by both direct and indirect measures of memory.

EXPERIMENT 2: THE IMPACT OF TO-BE-FORGOTTEN INFORMATION ON INDIRECT MEASURES OF MEMORY[3]

The basic goal of Experiment 2 was to determine whether an instruction to forget inhibits overall access to the to-be-forgotten item in memory or, rather, inhibits access to those items because they are part of an episode that is inhibited—that is, the to-be-forgotten list. Additionally, the experiment was designed to determine whether to-be-forgotten items, though inhibited as measured by conscious recall and/or the absence of proactive interference owing to those items, might nonetheless continue to have indirect, or unconscious, effects on performance.

Method

Subjects. The subjects were 48 undergraduates attending the University of California, Los Angeles, who either volunteered to participate in the study or participated to fulfill a course requirement.

Design and procedure. Much of the basic design and procedural details of the present experiment are similar to those of Experiment 1. All subjects were presented with the three types of lists illustrated in Fig. 1: remember–remember, forget–remember, and blank–remember. As before, each type of list was followed equally often either by an immediate recall test of the to-be-remembered words or by a recall test delayed by an intervening task. However, the intervening task was a word-fragment-completion task, rather than a recognition memory task.

Because a word-fragment-completion task was employed in the present experiment, the words comprising the lists were changed from common, high-frequency words to lower-frequency words (e.g., *dolphin, liberty, ridicule, violin*) in order to make priming effects more apparent. Because of the increased difficulty of these words, only 12 words were presented before and after the mid-list cues to remember or to forget, which were again rows of plus signs and minus signs, respectively. In addition, each word was presented for 5 s rather than 3 s, as were the mid-list forget and remember cues. For the blank–remember lists, subjects again performed a similar judgment task with respect to pairs of shapes. One minor procedural difference was

[3] Portions of Experiment 2 were reported in a paper by Bjork, Bjork, and Kilpatrick (1990, November) presented at the meeting of the Psychonomic Society (New Orleans, LA).

that the word lists and the stimuli for the shape-discrimination task were presented on a computer screen rather than via a slide projector.

Each subject received each type of list twice, with each type being tested once in each of the two types of testing conditions: a recall test of all to-be-remembered words that was either immediate or delayed by an intervening word-fragment-completion task. For all list types, the word-fragment-completion tasks contained 12 words of which 6 were always new words and 3 were always postcue to-be-remembered words from the current list. For remember–remember and forget–remember lists, the 3 remaining words on the fragment-completion task were taken from the precue part of the current list—that is, they were 3 precue to-be-remembered words or 3 to-be-forgotten words, respectively. For blank–remember lists, the remaining 3 words were additional new words. The 3 precue and postcue words appearing on the word-fragment-completion task were randomly selected from each list with the constraint that only 1 word could be taken from roughly each third of the appropriate list half. Similarly, the position in which they appeared in the word-fragment-completion task itself was randomly determined with the constraint that 1 precue and postcue word would appear in roughly each third of the task.

As before, booklets were handed out before the start of each list, and subjects could not tell from these booklets which type of test would be given for the current list. During the word-fragment-completion task, the word fragments were presented one at a time for 10 s on the computer screen, and subjects wrote their responses in the booklet, but in such a way that they could not see any prior responses they had made. Subjects had been instructed that although some of the words in the word-fragment-completion tasks might be ones that had previously been presented, any word that correctly completed the fragment was a correct response.

Finally, as with Experiment 1, counterbalancing procedures were used such that, across subjects, each specific list of words occurred equally often as each type of list, each specific list half occurred equally often as the first and second half of a total list, and each combination of list type and testing condition appeared equally often in all presentation positions. Thus, no item differences or presentation-order differences are involved when making comparisons across list type conditions, testing conditions, or between to-be-forgotten and to-be-remembered words. Additionally, counterbalancing procedures were used such that the precue and postcue words appearing on the word-fragment-completion task for half the subjects appeared as new words for the other half of the subjects and vice versa. Thus, no item differences need to be considered when testing for priming effects on the word-fragment-completion task—that is, when comparing differences in the completion rates for old versus new words.

Results

Immediate recall test. The correct recall percentages obtained for postcue to-be-remembered words that had *not* appeared on the word-fragment-completion task are shown in Fig. 4 for the three list types when the recall test was immediate. Once again, the basic pattern obtained in the immediate recall condition of Experiment 1 and the Bjork et al. (1973) study appeared in the present study as well. Planned-

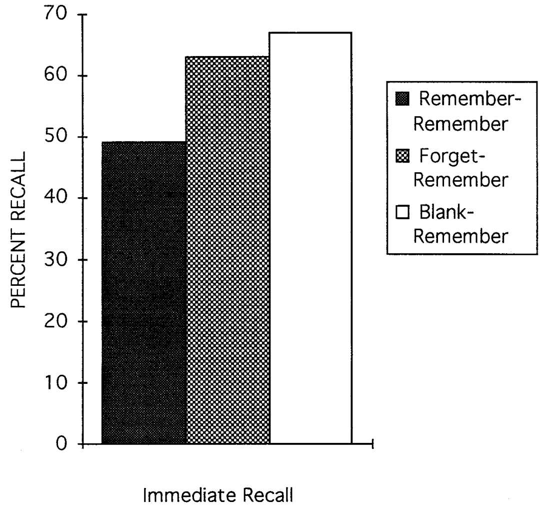

FIG. 4. Percentages of correct recall for postcue to-be-remembered words that had *not* appeared on the word-fragment-completion task for each list type on the immediate recall test.

comparison tests confirmed this pattern: Recall of postcue to-be-remembered words was significantly lower in the remember–remember list condition than in the forget–remember condition, and the difference in recall performance between the forget–remember and blank–remember conditions was not significant, $F(1, 47) = 24.54$, $MS_e = 249.46$, $p < .0001$; $F(1, 47) = 1.40$, $MS_e = 206.61$, $p > .24$.

Word-fragment-completion task. Given that to-be-forgotten items were inhibited, as evidenced by their lack of interference on the recall of to-be-remembered items, two critical questions remain. First, was access to the to-be-forgotten items—as appropriate completions on the word-fragment-completion task—also inhibited? That is, was the completion rate for precue to-be-forgotten words significantly lower than that for precue to-be-remembered words and, possibly, not different from the completion rate for new, or unprimed, words? A positive answer to this question would indicate that the effects of the inhibitory processes initiated by the cue to forget were not limited to inhibiting conscious access to the precue list-learning episode, but also extended to the inhibition of specific-item representations in semantic memory—that is, that the activation of the specific to-be-forgotten items themselves has been inhibited. The results shown in Table 1 contain the answer to this first question.

Table 1 presents the proportions of word fragments that were completed on the word-fragment-completion task with words that came from the first (precue) or second (postcue) half of each list or with the new words that corresponded to them, referred to as New-Word Targets in the table. (Across subjects, given the counterbalancing procedures in Experiment 2, the precue, postcue, and new words that were the targets of the word-fragment task were the same pool of words.) The result of primary interest—as shown in Table 1—is that there was no significant difference in the extent to which to-be-forgotten and to-be-remembered precue (list 1) words primed performance on the fragment-completion task, as evidenced by the absence of a significant interaction between list type (remember–remember vs. forget–

TABLE 1
Proportions of Word Fragments Completed with Precue-, Postcue-, and New-Word Targets for Each List Type (Priming Shown in Parentheses)

	List type		
Wordtype	Remember–remember	Forget–remember	Blank–remember
New: Unprimed	.27	.25	.27
Precue: First list half	.81	.74	—
(precue − new)	(+.54)	(+.49)	(—)
Postcue: Second list half	.78	.78	.75
(postcue − new)	(+.51)	(+.53)	(+.48)

remember) and word type (precue vs. new), $F(1, 47) = 2.03$, $MS_e = .43$, $p > .16$. Additionally, there was no significant difference in the extent to which a word's presentation in the postcue list following an instruction to forget or to remember increased the likelihood of its being supplied as a completion for the appropriate word fragment, as also indicated by the lack of a significant interaction between list type (remember–remember vs. forget–remember) and word type (postcue vs. new), $F(1, 47) = 0.58$, $MS_e = .73$, $p > .45$. Furthermore, the overall completion rates for precue and postcue words did not differ significantly, nor did the completion rates for precue and postcue words interact with the type of cue (remember vs. forget), $F(1, 47) = 0.98$, $MS_e = .53$, $p > .33$; $F(1, 47) = 1.63$, $MS_e = .32$, $p > .21$; respectively. Finally, considering only postcue (second list) words, priming rates did not differ across all three list types, $F(2, 94) = 0.50$, $MS_e = .64$, $p > .50$.

Thus, with respect to the first question posed above, it appears that the to-be-forgotten words primed word-fragment completion to roughly the same extent as corresponding to-be-remembered words. That is, subjects' access to to-be-forgotten words as lexical entries—in response to a word-fragment cue—was either not inhibited by the directed-forgetting instructions or, possibly, any such inhibition was released during the word-fragment-completion task, analogous to recognition testing that involves reexposure to List-1 items. The second critical question, then, is whether the intervening word-fragment-completion task reinstated the proactive interference owing to to-be-forgotten items. The answer to this question is contained in the results displayed in Fig. 5, which presents the delayed recall performance obtained for each type of list.

Delayed recall test. The percentages of correct recall shown separately for each list type in Fig. 5 are based on the recall of only those postcue to-be-remembered words that had *not* appeared on the word-fragment-completion task. Comparing the results for the delayed recall test shown in Fig. 5 to those for the immediate recall test shown in Fig. 4, it is clear that the overall level of recall performance has been reduced. The overall pattern of recall performance, however, appears to be the same as it was in the immediate recall condition. That is, there has apparently been no reinstatement of the proactive interference of the to-be-forgotten items. This apparent pattern was confirmed by the results of planned comparison tests, which revealed recall performance on remember–remember lists to be significantly poorer than recall

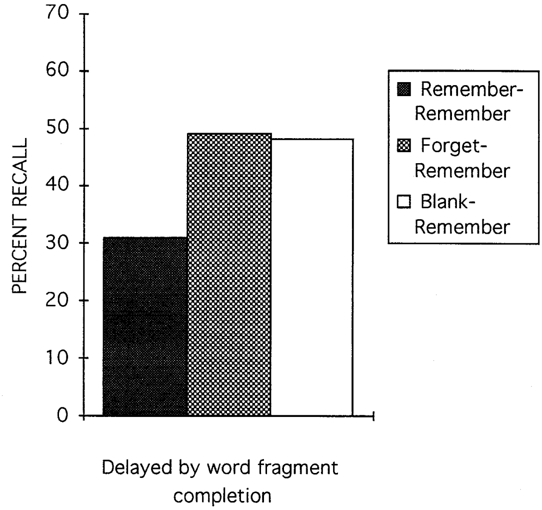

FIG. 5. Percentages of correct recall for postcue to-be-remembered words that had *not* appeared on the word-fragment-completion task for each list type on the delayed recall test.

performance on forget–remember lists, whereas the latter did not differ significantly from the recall performance obtained for blank–remember lists, $F(1, 47) = 6.01$, $MS_e = 377.28$, $p < .018$; $F(1, 47) = 1.47$, $MS_e = 332.42$, $p > .23$. Thus, while there was no sign that to-be-forgotten items were inhibited on the word-fragment-completion task (priming effects were just as strong for them as for the to-be-remembered words), that they are inhibited in some way is indicated by the lack of any proactive interference effects on the delayed recall test for forget–remember lists.

Discussion

The combined results of the present series of studies reveal some interesting properties of the retrieval inhibition evoked by directed-forgetting instructions. Clearly, special conditions appear to be necessary to reinstate the proactive interference owing to the to-be-forgotten items—that is, to release the inhibition that has been imposed on such items in response to the instruction to forget them. One such condition is that at least some subset of the to-be-forgotten items must be reexposed to subjects. Mere exposure, however, is not sufficient. During this exposure, the forgotten material must be processed in a manner that accesses, or makes contact with, the initial learning episode.

This latter necessary condition for inhibition release would seem to indicate something about the nature of the inhibition evoked in the directed-forgetting situation. It would not appear to be a general inhibition of the to-be-forgotten items as lexical entries—otherwise, there should have been less priming for to-be-forgotten than for to-be-remembered items on the word-fragment-completion task. Whereas the word-fragment-completion task does involve a type of retrieval—only some letters of each word were presented—the type of retrieval involved is largely data driven. It is not a task that directs or refers the subject back to the initial learning event or ‘‘target’’

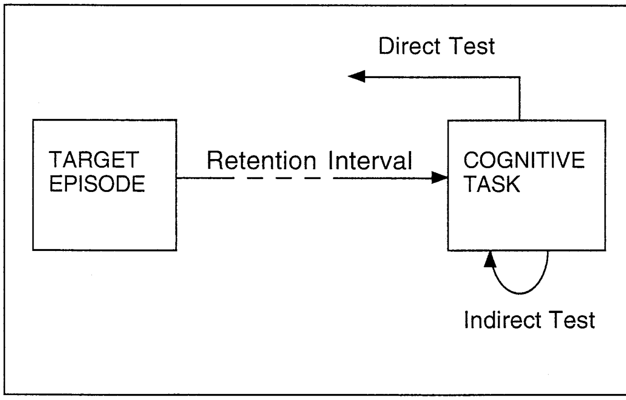

FIG. 6. Schematic representation of the different relationships between direct and indirect tests of memory and the target episode.

episode. The intervening recognition test, on the other hand, is just such a task, and, indeed, when subjects encountered to-be-forgotten items in the context of this type of task, proactive interference owing to the forgotten items was reinstated. Based on the pattern of results observed across the three present studies, then, the inhibition involved in the directed-forgetting situation appears to be a type of retrieval inhibition that impairs conscious access to original learning episodes that are the object of a forget instruction—that is, the episode in which the information was first learned and then intentionally forgotten. In contrast, the inhibition initiated by the instruction to forget does not seem to inhibit the activation level of to-be-forgotten information or to prevent it from having indirect or unconscious influences on behavior.

GENERAL DISCUSSION

We began this paper by commenting on the variety of situations in which we encounter either explicit or implicit cues to forget, and we stressed the adaptive aspects of forgetting in terms of the needs we have to update our memories. We also stressed that retrieval inhibition is a critical aspect of such forgetting. Under some conditions, however, as we have demonstrated, interference owing to the previously inhibited out-of-date information can return full blown. Additionally, even when to-be-forgotten information remains inhibited, in the sense that conscious access to that information is impeded—and proactive interference owing to that information is reduced, such information can continue to influence our judgments and behavior. In the remainder of this section, we speculate on what the properties of the retrieval inhibition evoked by cues to forget imply with respect to our ability to perform optimally in different types of situations. To aid this discussion, we first interpret the findings of the present studies in terms of the distinction between direct and indirect measures of memory.

The word-fragment-completion task employed in Experiment 2 of the present paper is an *indirect* test of memory as defined by Johnson and Hasher (1987) and by Richardson-Klavehn and Bjork (1988). As illustrated in Fig. 6, such tests or cognitive tasks are indirect in that they measure a subject's memory for a prior episode by assessing the impact of that episode on performance in a task that does not refer to the target episode. Thus, for example, if the prior presentation of a given word increases the likelihood that a subject can complete that word from letter fragments (R__ __S__N), or from a stem (REA______) or ending (______SON), or can identify that word when it is briefly exposed, we infer that the prior episode involving that word had an impact on the subject's memory. Examples of other indirect tests are perceptual identification, homophone biasing, evaluative responses, and certain nonverbal changes. *Direct* tests, such as tests of recall and recognition, on the other hand, instruct the subject to refer directly to the target episode itself.

In terms of the direct–indirect distinction, then, a way of looking at the results of Experiment 1 is to say that when to-be-forgotten items were encountered in the context of a direct test of memory—that is, one that referred subjects back to the learning that was the object of a forget instruction—the inhibition of those items was released. Interpreted in this way, it is rather easy to see that these results have important implications for how certain types of work situations or contexts should be structured—or perhaps, more to the point, should not be structured. For example, in occupations or job contexts that involve updating and/or frequent changing of plans—occupations that could range from short-order cooks to air-traffic controllers to any number of white-collar jobs where current accounts or cases involve different values of the same variables involved in prior cases—any reexposure to items that have been forgotten, in response either to a direct or to an implicit cue to forget, should not also ask for a judgment of some kind that makes reference to the out-of-date context. The use in certain work situations of a checkoff procedure in which now out-of-date items appear, for example, may be counterproductive if, as a result, such information is then more likely to intrude into the new context where it is no longer appropriate and will be a source of confusion.

Similarly, if we look at the results of Experiment 2 in terms of this distinction, we could say that while the direct retrieval of to-be-forgotten items was impaired, their indirect retrieval (on the word-fragment-completion task) was not. Consistent with this way of interpreting the results of Experiment 2, Basden, Basden, and Gargano (1993) recently reported obtaining equivalent priming of to-be-forgotten and to-be-remembered words on indirect tests of memory, including word-fragment-completion, for their studies employing the list method of directed forgetting. Thus, the present results as well of those of Basden et al. raise the issue of whether other analogous dissociations exist—that is, whether there are other instances in which subjects' intentional recall of to-be-forgotten items is impaired, but the indirect influence of those items on some other task remains unabated.

A likely case where such dissociations may exist is in the domain of impression formation. There is now a substantial body of literature (e.g., Carretta & Moreland, 1983; Golding, Fowler, Long, & Latta, 1990; Thompson, Fong, & Rosenhan, 1981; and Wyer & Budesheim, 1987; see also Johnson, 1994, for a review) showing that to-be-disregarded information continues to influence subjects' evaluative judgments. This research, often carried out in a simulated legal setting, tends to show that information ruled inadmissible on some basis remains in memory at essentially full strength as measured by its influence on later evaluative impressions of people, or on judgments of guilt or innocence.

We conclude with the speculation that, under certain circumstances, to-be-forgotten information may have a *larger* influence than would corresponding to-be-remembered information. Jacoby and his co-workers (see, e.g., Jacoby & Kelley, 1987; and Jacoby, Kelley, & Dywan, 1989) have provided convincing evidence that we are often unaware of the influence of some prior experience on a current judgment of some type. In one series of experiments, for example, in which subjects are asked to judge whether each of a series of names (e.g., Sabastian Weisdorf) is famous or not—that is, whether a given name is the name of a celebrity of some kind—the prior presentation of a given nonfamous name in an earlier phase of the experiment increased the likelihood that subjects falsely judged that name to be famous, even though—24 h earlier—those names were clearly labeled nonfamous. That is, the familiarity of such names owing to their having been presented in an earlier phase of the experiment is apparently misattributed to real world exposures of one kind or another. Jacoby, Woloshyn, and Kelley (1989) have shown further that such misattributions are more frequent when the initial presentation of such nonfamous names takes place under conditions of divided attention. Under such conditions, subjects are less able to identify the earlier experimental episode as the source of the familiarity of a given name. In general, such results are consistent with the well-known ‘‘sleeper’’ effect in social psychology: that the influence of a noncredible source of information can grow over time—that is, can grow as subjects lose the ability to attribute their memory for that information to an unreliable or untrustworthy source.

The research of Jacoby and his colleagues, taken together with the evidence that instructions to forget do not reduce appreciably the impact of the to-be-forgotten information on subsequent indirect measures of memory, has an interesting implication. If subjects suffer a kind of source amnesia for to-be-forgotten information, analogous to that fostered by divided attention, the unconscious influence of the information may actually be greater than the influence of comparable to-be-remembered information. That is, in a judgment task where it is important to identify a to-be-remembered or to-be-forgotten episode as the source of potentially misleading information, subjects may be less able, in the case of to-be-forgotten information, to identify the episode in question as the source of that information.

As indicated in the introduction, that subjects may be poor at identifying a to-be-forgotten episode as the source of information is suggested by the results of research by Evans and Kihlstrom (1973; see also Kihlstrom, 1983) on posthypnotic amnesia and by the results of research by Geiselman et al. (1983) on directed forgetting with subjects in the normal waking state. In both situations, when to-be-forgotten items *are* recalled, subjects are confused as to the source of those items compared to the recall of comparable to-be-remembered items, and the order of recall of such items is disrupted as well.

Considering all of these results as well as those obtained in the present studies, the possibility that information we have intentionally tried to forget could have greater unconscious influences on our thoughts, actions, and judgments than comparable to-be-remembered information seems not entirely unlikely. And, if not so unlikely, this possibility has both interesting and alarming implications, as in the courtroom setting described above and in similar settings, where jurors and other people may be unable consciously to recall misleading or prejudicial information that they were instructed to disregard, yet they may continue to be influenced by it in ways of which they are unaware.

REFERENCES

Basden, B. H., Basden, D. R., & Gargano, G. J. (1993). Directed forgetting in implicit and explicit memory tests: A comparison of methods. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **19,** 603–616.

Bjork, E. L., Bjork, R. A., & Glenberg, A. M. (1973, November). *The reinstatement of interference owing to to-be-forgotten items.* Paper presented at the meeting of the Psychonomic Society, St. Louis, MO.

Bjork, E. L., Bjork, R. A., & Kilpatrick, H. A. (1990, November). *Direct and indirect measures of inhibition in directed forgetting.* Paper presented at the meeting of the Psychonomic Society, New Orleans, LA.

Bjork, E. L., Bjork, R. A., & White, S. (1984, November). *On the induced recovery of proactive interference.* Paper presented at the meeting of the Psychonomic Society, San Antonio, TX.

Bjork, R. A. (1970). Positive forgetting: The noninterference of items intentionally forgotten. *Journal of Verbal Learning and Verbal Behavior,* **9,** 255–268.

Bjork, R. A. (1972). Theoretical implications of directed forgetting. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory.* Washington, DC.: Winston & Sons.

Bjork, R. A. (1978). The updating of human memory. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 12, pp. 235–259). New York: Academic Press.

Bjork, R. A. (1989). Retrieval inhibition as an adaptive mechanism in human memory. In H. L. Roediger III & F. I. M. Craik (Eds.), *Varieties of memory & consciousness: Essays in honour of Endel Tulving* (pp. 309–330). Hillsdale, NJ: Erlbaum.

Block, R. A. (1971). Effects of instructions to forget in short-term memory. *Journal of Experimental Psychology,* **89,** 1–9.

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

Elmes, D. G., Adams, C., & Roediger, H. L. (1970). Cued forgetting in short-term memory: Response selection. *Journal of Experimental Psychology,* **86,** 103–107.

Evans, F. J., & Kihlstrom, J. (1973). Posthypnotic amnesia as disrupted retrieval. *Journal of Abnormal Psychology,* **82,** 317–323.

Geiselman, R. E., & Bagheri, B. (1985). Repetition effects in directed forgetting: Evidence for retrieval inhibition. *Memory & Cognition,* **13,** 57–62.

Geiselman, R. E., Bjork, R. A., & Fishman, D. (1983). Disrupted retrieval in directed forgetting: A link with posthypnotic amnesia. *Journal of Experimental Psychology: General,* **112,** 58–72.

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

Gross, A. E., Baresi, J., & Smith, E. E. (1970). Voluntary forgetting of a shared memory load. *Psychonomic Science,* **20,** 73–75.

Jacoby, L. L., & Kelley, C. M. (1987). Unconscious influences of memory for a prior event. *Personality and Social Psychology Bulletin,* **13,** 314–346.

Jacoby, L. L., & Kelley, C. M., & Dywan, J. (1989). Memory attributions. In H. L. Roediger III & F. I. M. Craik (Eds.), *Varieties of memory & consciousness: Essays in honour of Endel Tulving* (pp. 391–422). Hillsdale, NJ: Erlbaum.

Jacoby, L. L., Woloshyn, V., & Kelley, C. M. (1989). Becoming famous without being recognized:

Unconscious influences of memory produced by dividing attention. *Journal of Experimental Psychology: General,* **118,** 115–125.

Johnson, H. M. (1994). Processes of successful intentional forgetting. *Psychological Bulletin,* **116,** 274–292.

Johnson, M. K., & Hasher, L. (1987). Human learning and memory. *Annual Review of Psychology,* **38,** 631–638.

Kihlstrom, J. F. (1983). Instructed forgetting: Hypnotic and nonhypnotic. *Journal of Experimental Psychology: General,* **112,** 73–79.

Reed, H. (1970). Studies of the interference processes in short-term memory. *Journal of Experimental Psychology,* **84,** 452–457.

Reitman, W., Malin, J. T., Bjork, R. A., & Higman, B. (1973). Strategy Control and directed forgetting. *Journal of Verbal Learning and Verbal Behavior,* **3,** 244–254.

Richardson-Klavehn, A., & Bjork, R. A. (1989). Measures of memory. *Annual Review of Psychology,* **39,** 475–543.

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

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

Received December 12, 1995