Interference and Inhibition in Memory Retrieval

Michael C. Anderson
James H. Neely

I. INTRODUCTION

In memory research, interference refers to the impaired ability to remember an item when it is similar to other items stored in memory. Consider, for example, the deceptively simple task of recalling where you parked your car at a local shopping center. If you have never before been to that shopping center, recalling your car's location may be fairly easy. If you park there frequently, however, you may find yourself reunited with the spot where you parked yesterday or, if you are like the present authors, standing befuddled at the lot's edge. Further, if asked where you parked on previous visits, you would almost certainly fail to recall the locations, as though your intervening parking experiences had overwritten those aspects of your past. These examples illustrate typical cases of retrieval failure arising from interference. Understanding the causes of such interference has been a central goal of research on forgetting since the inception of experimental psychology.

Research on interference has generated a variety of conceptions of how forgetting occurs. This variety may be illustrated intuitively in terms of our previous parking situation. For example, sometimes our ability to recall our current parking location seems blocked by the intrusion of similar episodes. When this occurs, we often feel confident that we know where we parked, but that recall of the location demands that we penetrate through memories

Memory

that get in the way. On other occasions, particularly when retrieving older parking episodes, we feel less confident that we retained the information, as though subsequent episodes disrupted or overwrote prior experience. Additional possibilities are suggested by different examples. For instance, we often experience difficulty recalling our new telephone number after a recent change in residence, because our old number keeps coming to mind, disrupting recall of the new one. With time, however, we typically suppress the outdated number. In terms of our earlier parking example, this intuition suggests that ensuring the recallability of today's parking spot involves the suppression of earlier parking memories in that same lot. Each of these intuitions has motivated theorizing at some point in the history of interference research.

The present chapter reviews what experimental research has revealed about the causes of memory interference and the breadth of situations in which these mechanisms operate. We describe this research in four main sections. First, we discuss some widely held assumptions about the situation of interference, focusing on the idea that such effects arise from competition for access via a shared retrieval cue. This notion is sufficiently general that it may be applied in a variety of interference settings, which we illustrate briefly. Having introduced these basic assumptions, we review the classical interference paradigms from which these ideas emerged, as well as the variety of particular conceptions of forgetting developed in the context of these procedures. Many of these ideas remain relevant today, influencing how we conceive interference in modern terms. In the next section, we move outside of the classical arena to review more recent phenomena that both support and challenge classical conceptions of interference. These phenomena provide compelling illustrations of the generality of interference and, consequently, of the importance of our understanding its mechanisms. We close by highlighting a recent perspective on interference that builds upon insights from modern work, while validating intuitions underlying several of the classical interference mechanisms. According to this new perspective, forgetting derives not from acquiring new memories per se, but from the impact of later retrievals of the newly learned material. After discussing findings from several paradigms that support this retrieval-based view, we illustrate how forgetting might be linked to inhibitory processes underlying selective attention.

II. BASIC ASSUMPTIONS OF INTERFERENCE RESEARCH

Most approaches to interference share basic assumptions about the representations and retrieval processes at work in interference situations. For example, many approaches characterize retrieval as a progression from one or more retrieval cues to items stored in memory by way of associative links. Retrieval cues can be anything from components of the desired memory to incidental concepts associated with that item during its encoding. Thus, recalling where we parked our car might involve the activation of many concepts, including features of the car, the act of parking, the layout of the lot, or the time of day, any of which might reasonably have been encoded when we parked. The success of this progression from cues to the target memory hinges on many factors, including the number of cues used and the strengths of the associations linking cues to the memory items, both of which are influenced by the amount and character of attention paid during encoding. Under normal circumstances, retrieving a target item is thought to occur when the cues available at the time of recall are sufficiently related to that target to identify it uniquely in memory.

By the foregoing analysis, the essential problem in interference is that the retrieval cues available at the time of recall fail to access the target memory. Why might such failures occur? Figure 1 illustrates one general approach to this question. In this approach, interference arises when the retrieval cue normally used to access a target (Figure 1A) becomes associated to additional memory items (Figure 1B). Successfully progressing from a retrieval cue to a target memory thus depends not only on how strongly that cue is related to the target, but also on whether the cue is related to other items in memory as well. When a cue is linked to more than one item in memory, those items are assumed to compete with the target for access to conscious awareness—what M. C. Anderson, Bjork, and Bjork (1994) have referred to as the *competition assumption*. In the present chapter, we refer to any negative effect on memory performance associated with this competition as *interference*. Interference owing to the mechanisms of competition is generally thought to increase with the number of competitors, a notion supported by the observed tendency for recall performance to decrease with the number of items that are paired with the same cue. This generalization has come to be known as the *cue-overload principle* (see, e.g., Watkins, 1978).

How might the notion of competition among items associated to the same retrieval cue capture interference arising from the acquisition of new memories similar to those already stored? The basic approach to this question can be illustrated in terms of our previous parking example. During your visit to the shopping center, you encoded aspects of your parking experience into a mental representation of that event. Other parking experiences that are similar to this visit will also contain characteristics that are stored in the target event, including, for instance, the fact that you drove the car, the type of car you drove (e.g., a 1989 blue Honda), and perhaps your goal of doing shopping at the supermarket. If components of the target event (e.g., your concepts of yourself, parking, and your Honda) serve as the primary retrieval cues by which you access your car's location, other memories sharing those features will also be evoked during retrieval. Figure

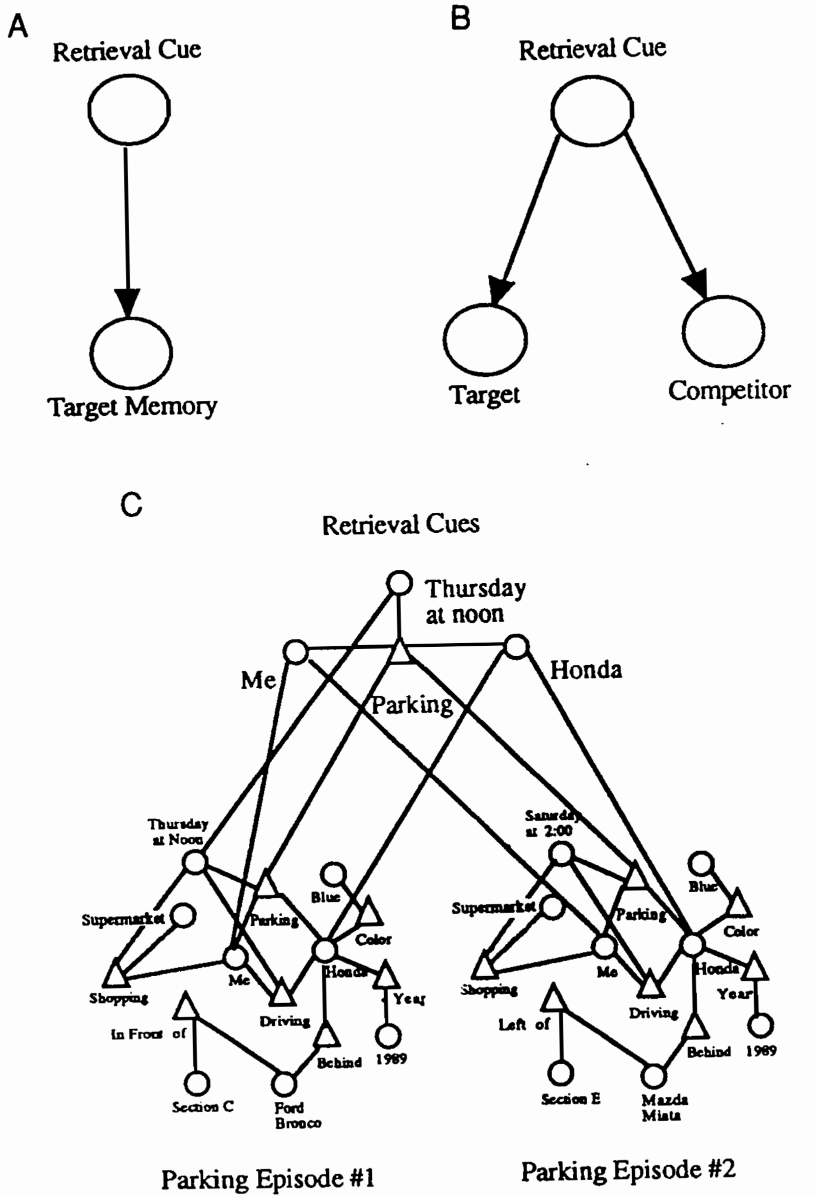

FIGURE 1 Illustration of the notion of competition among items sharing the same retrieval cue. (A) A retrieval cue that is associated to only one target item in memory. (B) The basic situation of interference, in which a retrieval cue becomes associated to one or more additional competitors that impede recall of the target, given presentation of the shared retrieval cue. (C) How the basic situation of interference illustrated in (B) may be applied to understand a more complex example of interference in which two episodes of having parked at the supermarket interfere because they share the retrieval cues "Me," "Honda," and "Parking" at the time of retrieval. Circles and triangles in the representations of Episodes 1 and 2 in (C) depict concepts and relations, respectively.

1C illustrates this point by showing how the general logic advanced in the case of one cue (Figure 1B) may be scaled up to the many cues available in the more naturalistic example of remembering where one parked (e.g., "Me," "Parking," and "Honda"). If retrieval cues become less effective as they acquire new associations, then similar memories should compete with one another to the extent that the sets of cues useful in accessing those memories overlap. Thus, competition among items that share the same retrieval cues provides a useful way of viewing the problem of interference, even for complex episodes.

The notion of competition among items that share retrieval cues is even more general than the previous example. For instance, items in memory need not be episodes to compete with one another. Consider the task of retrieving, from general knowledge, the particular exemplar *banana* given the cue *fruit*. This simple retrieval should be impeded by competition from other fruits that one might know, such as *orange* and *lemon*. Furthermore, items need not be semantically similar to compete, provided that they share a common cue. For instance, competition might ensue among items grouped by mere co-occurrence in time, as in the problem of recollecting, from the morning conversation with one's spouse, the items one is supposed to buy at the market. Although the items to be remembered may be quite dissimilar, they might compete with one another if the memory for the morning episode serves as the primary cue. A recurring theme of this chapter is the far-ranging generality of retrieval competition as a mechanism for interference.

III. CLASSICAL APPROACHES TO INTERFERENCE

Much of modern thinking about the causes of interference has been shaped by the substantial body of research completed during what has come to be known as the "classical interference era" (approximately 1900–1970). Indeed, the basic insight that interference arises from the competition of memory items for a shared retrieval cue originated early in this period (McGeoch, 1936, 1942). In this section, we review the basic paradigms of classical interference research because these paradigms provide a clear specification of the conditions of interference. We then discuss the wealth of findings and theoretical concepts emerging from these paradigms. These ideas remain relevant to contemporary theorizing and to our intuitions about the causes of forgetting. (The reader is referred to Postman, 1971; Postman & Underwood, 1973; Crowder, 1976, for more detailed reviews of the classical interference literature.)

A. Methodology of Classical Interference Research

If interference arises from competition for a shared retrieval cue, an experimental paradigm for studying this phenomenon should allow control over three things: (1) the cues to which people associate target memories; (2) the cues by which people retrieve the target at the time of test; and (3) the relations of those cues and targets to other items in memory. Control over these factors has been achieved in the classical paired-associate paradigm. We now describe this paradigm and the basic methodologies and designs of classical interference experiments.

1. The Paired-Associate Paradigm

According to the classical interference perspective, the participant's task in the paired-associate paradigm (and in learning more generally) is to acquire memory *responses* to verbal *stimuli*. In this paradigm, a person studies unrelated pairs of verbal items (typically words), one pair at a time, for a later memory test. For example, people might be shown the pair *dog–rock* and be instructed to study this pair so that when the stimulus term *dog* is later provided as a cue, it can elicit the relevant response term, *rock*. Typically, repeated study–test trials would be given until participants achieved a certain criterion level of learning on these responses (a practice not generally observed in modern research using paired associates). Learning paired associates in this fashion is typically thought to induce the encoding of various associations between the verbal items, including forward associations linking stimuli to responses (i.e., the association linking *dog* to *rock*), backward associations linking the responses back to their stimuli (i.e., the association linking *rock* to *dog*), and contextual associations linking each item to the general representation of the list context. Most of our discussion focuses on forward associations, although other kinds of associations are considered when relevant.

Interference is studied in the paired-associate paradigm by having people learn a first list of critical target pairs (e.g., *dog–rock*) and then a second list of pairs that bear any one of a number of different relations to the critical target pairs, with the aim of examining the effect of such new learning on recall performance for the initial target item (e.g., *rock*). The relation between the new paired associates and those previously learned are described in a standard notation in which a cue–target pair from the first list (e.g., *dog–rock*) is designated as an A–B item, with the first letter of this notation referring to the stimulus term and the second letter to the response. Following this notation, the items on the second list can be related to those on the first list in a variety of ways, four of which are illustrated in Figure 2A. Note that all four groups of participants depicted in this example study the same initial A–B list (List 1), but the content of the second list they study (List 2) varies. In the A–B, A–B paradigm (the leftmost group), the pairs on the second list (A–B) share both the stimulus and the response terms of the first-list (A–B) pairs (e.g., *dog–rock*) and are thus merely repetitions of those

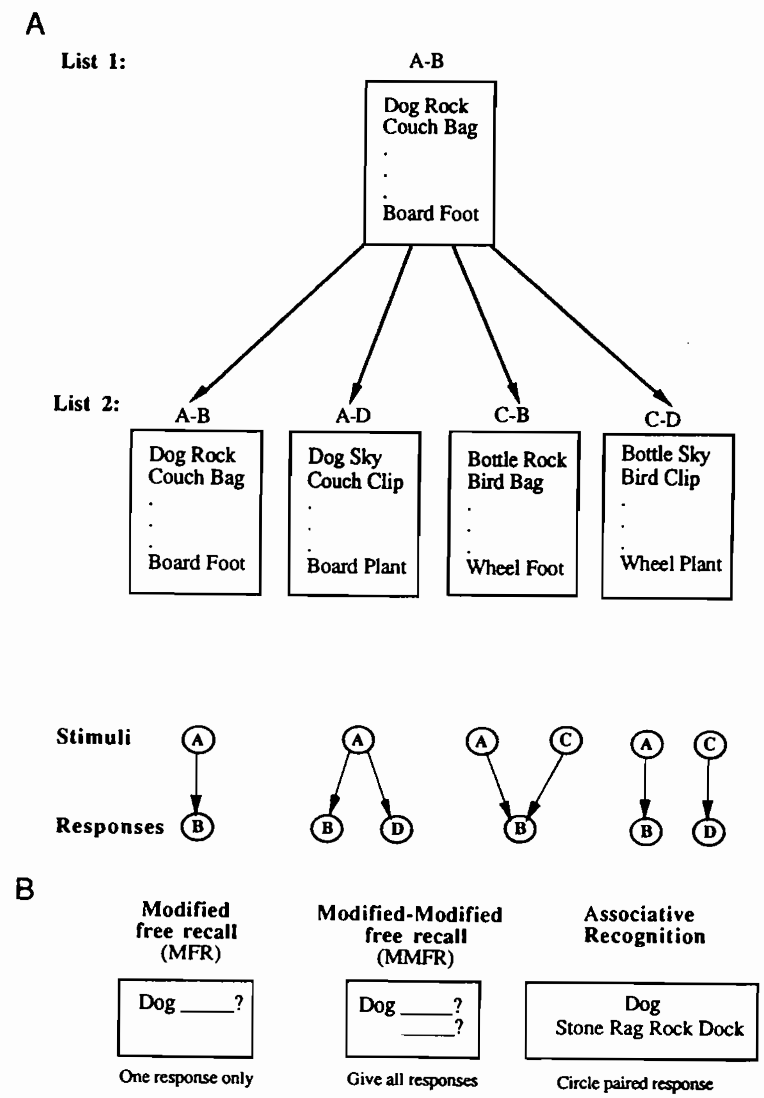

FIGURE 2 Classical interference methodology. (A) The four most common interference designs used in classical interference research, in which an initial A–B list (List 1) is followed by another list (List 2) that presents either A–B, A–D, C–B, or C–D associates (see text for elaboration). The boxes in this figure do not imply that items were presented together on a page; they denote only that the items appeared on the same list of individually presented associates. The stimulus–response relations resulting from such consecutive presentations are depicted schematically below each condition, with the forward associations between the various retrieval cues and memory responses represented. (B) Three common testing procedures employed in paired-associate research.

pairs; in the A–B, A–D paradigm, the second-list (A–D) pairs share the stimuli of the first-list (A–B) pairs but differ in their responses (e.g., *dog–sky*); in the A–B, C–B paradigm, the second list (C–B) pairs share the response terms of the first-list (A–B) pairs, but differ in their stimuli (e.g., *bottle–rock*); finally, in the A–B, C–D paradigm, the second-list (C–D) pairs share neither the stimuli nor the response terms with the first-list (A–B) pairs (e.g., *bottle–sky*). These critical relationships between the first-list and second-list associates are depicted schematically in Figure 2A, immediately beneath the groups studying those relationships. (Note that only forward associations are depicted.) Other relationships may also be constructed by manipulating stimulus or response similarity, introducing additional stimuli or responses, and re-pairing the same stimuli and responses in different ways (Postman, 1971). For present purposes, the most important relationship is found in the A–B, A–D design. This design represents the classical situation of interference in which memory targets (B and D) share a retrieval cue (A) that induces competition when presented as the cue on later tests (i.e., *rock* and *sky* should compete when *dog* is presented).

After two lists of paired associates have been studied, memory for the target responses from one or both lists can be assessed in a variety of ways. Several of the most typical testing procedures are depicted in Figure 2B. The leftmost diagram in Figure 2B depicts the test, used in many studies conducted prior to 1959, in which the experimenter provided the stimulus for each pair and asked the subject to retrieve the first response that came to mind—a procedure called *modified free recall* (Underwood, 1948; hereafter, MFR). This procedure was a change from earlier recall tests in which participants were *directed* to retrieve one response from a particular list (e.g., List 1). To the right of the MFR procedure is depicted the test used in most studies conducted after 1959—the *modified modified free-recall* (hereafter, MMFR) procedure (Barnes & Underwood, 1959)—in which the experimenter asked for all responses associated with the stimulus and allowed participants as much time to recall as needed. Finally, the experimenter can also test a person's ability to recognize having seen a particular target response paired with a stimulus by providing the stimulus and a set of items, including the response, and asking the person to match the appropriate response paired with that stimulus (illustrated in the rightmost diagram). Although such associative recognition tests were explored late in the history of interference research, they have proven important as a means of discriminating among various theories of interference. Performance on all of these tests can be assessed by a number of dependent measures, but here we focus on the percentage of requested responses that are correctly recalled or recognized. (See Postman, 1971, for a discussion of other measures such as relative recall.)

2. Interference Designs

The paired-associate methodology just described can be used in numerous ways to study interference phenomena. Here, we describe the designs used to examine two of the most important and widely studied phenomena in classical interference research: retroactive interference and proactive interference. Findings from these paradigms have proven informative with respect to the mechanisms of forgetting.

a. Retroactive Interference

Retroactive interference refers to impaired memory performance on target items caused by learning new material between the initial encoding of those target items and their final test. For example, we suffer retroactive interference when we can no longer recollect where we parked last week because we have parked in the same lot on several subsequent occasions. This crucial phenomenon was the primary focus of classical interference theory for over six decades and was primarily studied using the classic retroactive interference design (G. E. Mueller & Pilzecker, 1900) illustrated in Figure 3A. In the experimental condition (left side of 3A), people study a first list of paired associates (upper box) and then a second list, the pairs of which may be

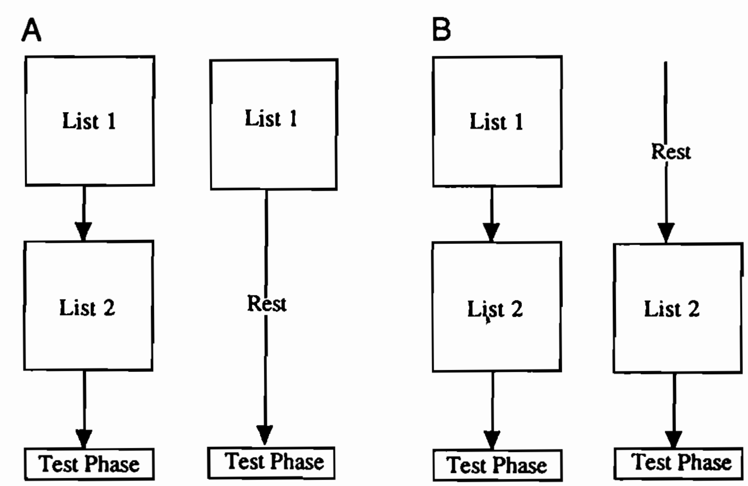

FIGURE 3 The generic procedure for the classical retroactive (A) and proactive (B) interference designs. The left portion of each panel depicts the experimental group, which always receives two lists of items, and is tested either on the first list (A) or on the second list (B). The right portion of each panel depicts the control group, which only receives one list, either at the same point as List 1 in the case of retroactive interference designs, or at the same point as List 2 for proactive interference designs. As noted in both panels, control participants rest (or perform irrelevant activity) in place of receiving an additional list.

related to first-list items in any of the ways discussed in our previous section on paired associate methodology (note that this procedure can also be used with stimuli other than paired associates). Following these two lists, people are tested with either an MFR, MMFR, or recognition test for the List 1 A–B responses. In the control condition, people also study a first list of responses, but either rest or engage in irrelevant activity during the interval in which people in the experimental condition study List 2. Thus, these conditions allow us to ask the crucial question "What is the effect of learning new information (i.e., List 2) on the ability to remember information that was previously studied (i.e., List 1), relative to a situation in which no additional information was learned at all (i.e., Control List 1)?"

A great number of studies using this paradigm have consistently demonstrated that memory performance on first-list A–B items suffers greatly when a second list of responses must be learned. The magnitude of this retroactive interference effect varies dramatically as a function of the stimulus conditions and the method of testing, with the greatest forgetting of A–B items found in the A–B, A–D paradigm with MFR tests. This generalization reinforces the notion that interference effects reflect competition among items sharing the same retrieval cue. However, some interference is often observed in the A–B, C–D procedure, as well (McGovern, 1964), although in dramatically reduced amounts. Interestingly, even when severe deficits in recall are demonstrated, our ability to recognize the older information often remains intact: recognition tests often show little or no retroactive interference in the A–B, A–D paradigm relative to the A–B, C–D paradigm (Postman & Stark, 1969). This finding suggests that recognition tests eliminate or greatly reduce the retrieval competition that occurs in the MFR and MMFR tests.

b. Proactive Interference

Proactive interference refers to previously learned materials hurting our memory for more recently learned items. For example, we suffer proactive interference when we fail to recall our new phone number momentarily because our old number intrudes during the recall process. Figure 3B depicts the design that allows us to examine proactive interference. The proactive interference paradigm resembles the retroactive interference design except that (1) it tests people's memory for the List 2 responses rather than the List 1 responses, and (2) in the control condition, the rest (or performance of irrelevant activity) replaces List 1 learning rather than List 2 learning. Thus, this design allows us to explore how previously acquired knowledge (i.e., List 1) might impair our ability to recollect new information (i.e., List 2), relative to a situation in which the previous knowledge had not been learned (Control, List 2).

Many studies using the proactive interference procedure have demonstrated that people are more likely to forget items from a list when a prior list has been studied. The magnitude of this proactive interference effect varies as a function of the stimulus structure and test conditions in the same way that retroactive interference does, with proactive interference being most severe when lists share retrieval cues. Furthermore, the effect of proactive interference is far more severe when recall is tested rather than recognition. In addition, the relative magnitudes of proactive and retroactive interference effects vary in interesting ways as the retention interval between List 2 and the final memory test increases. Whereas retroactive interference is more pronounced at short retention intervals, proactive interference dominates at longer delays (Postman, Stark, & Fraser, 1968), a finding that has proven important in arguments about the mechanisms of interference.

B. Classical Conceptions of Interference

The insight that interference might result from the competition of memories sharing the same cue provides only a starting point for understanding the causes of forgetting. Why might associating many memories to the same cue hinder our ability to recall items that were once accessible? As we noted in the introduction to this chapter, a variety of answers to this puzzle have been explored. We now review some of these ideas, highlighting both their classical origins and modern incarnations. This review is organized according to a general tripartite division of theories, within which retrieval failures are attributed to one of three components of the memory representations: (1) the retrieval routes (associations) linking cues to target memories; (2) the retrieval cues used to access targets; and (3) the activation levels in the memory targets themselves. (For a more detailed discussion of these and other theories, see M. C. Anderson & Bjork, 1994.) For the remainder of the present chapter, the more current terms *cue* and *target* are used synonymously with the classical terms stimulus and response, respectively.

1. Ineffective Retrieval Routes

Sometimes we discover that an item that we failed to recall on an earlier occasion was in memory all along. These failures to retrieve available information even occur when we use familiar methods to retrieve the information. For example, looking at the dial of a gym locker combination lock may fail to elicit the desired combination, no matter how many times one has opened the lock in the past, and no matter how intently one stares. In such cases, retrieval failure might derive from some breakdown in the retrieval method by which we normally access the desired information. In the aforementioned tripartite approach to interference, these intuitions suggest that the association between the cue and the target memory is degraded. This general approach is the oldest and most popular means of explaining forgetting. Theories that explain interference in terms of ineffective associations may be categorized in at least two different classes: occlusion and unlearning.

a. Occlusion

Failures to retrieve an item are sometimes accompanied by persistently intrusive memories that are similar to the target item. These situations are characterized by a distinctive feeling: Intrusive memories seem to get in the way or occlude the target item, and we feel confident that if we could only get past the intruding memory, we would access the target. The clearest examples of this experience occur during the tip-of-the-tongue state, in which our ability to name a particular person, place, or object seems thwarted by a persistently intrusive word. Such persistent alternates are typically similar to the target item on both semantic and linguistic dimensions, such as the competing words' lengths and initial sounds (R. Brown & McNeill, 1966; see A. S. Brown, 1991 for a review). The existence of situations analogous to occlusion in nonlanguage contexts—as when we sometimes visit yesterday's parking spot rather than today's—suggests that occlusion may cause interference more generally.

To understand how similar memories might occlude target retrieval, consider how such blocking might impair the recall of B items (i.e., produce retroactive interference) in the A–B, A–D paradigm. In this paradigm, presentation of the A cue should induce B and D to compete for access to conscious awareness, with the induced competition being greater the more strongly D is associated to A. When the D item is sufficiently strong, it may win this competition, displacing the retrieval of the B target item. Note that this view need not assume that anything has happened to the target memory B itself or to its association to the cue A. Rather, the *effectiveness* of the association linking A to B is presumed to decrease as the associations for competitors become more effective, inducing greater competition during retrieval. The notion that the amount of interference exerted by an item increases with its memory strength—which M. C. Anderson et al. (1994) have called the *strength-dependence assumption*—forms the cornerstone of classical and modern approaches to explaining interference phenomena (Gillund & Shiffrin, 1984; McGeoch, 1942; Mensink & Raaijmakers, 1988; Raaijmakers & Shiffrin, 1981).

The notion that competition is strength dependent—indeed, the more basic notion of competition for a shared retrieval cue itself—originates from McGeoch's (1942) classical analysis of the conditions of interference. According to McGeoch, interference is most strongly instantiated in the A–B, A–D paradigm, because in this paradigm people must acquire mutually

incompatible responses to a common retrieval cue. Responses thus related were assumed to impede each other's accessibility in a process that McGeoch called *response competition*. For example, when people were provided with an A stimulus in a List 1 recall test, both the B and D responses were assumed to compete with one another until one response attained momentary dominance and was reported. Because a List 1 recall test (and the subsequent MFR test) requested only one response for each cue, the ultimate resolution of this competition precluded participants from reporting the competing item. McGeoch assumed that the magnitude of this memory impairment or *reproductive inhibition* on a target response varied as a function of its competitor's strength of association to the shared cue. Thus, McGeoch's proposal was an early (perhaps the first) instance of an occlusion theory. (For clarity, in the remainder of the present chapter, we use the term response competition when referring to the general proposal that items sharing a cue compete, and occlusion when referring to models in which this competition is strength dependent and causes forgetting. Historically, however, McGeoch's theory of forgetting has been called the response-competition theory.) In support of these ideas regarding strength-dependent competition, McGeoch cited the tendency for the D responses to intrude on final memory tests for the B targets, and for retroactive interference to increase as competition from the D response was increased through greater amounts of A–D learning (but see the discussion of Melton & Irwin's, 1940, results later). Finally, this account provides a natural explanation for both proactive and retroactive interference because the tendency for items to compete with one another at the time of the final recall test should not depend on the order in which these items were initially learned.

The assumptions about competition and strength dependence proposed by McGeoch (1942) permeate modern mathematical and computational thinking with regard to the causes of interference. Clear examples of this influence can be found in the relative strength retrieval assumptions typically adopted in modern memory architectures (J. R. Anderson, 1983; Raaijmakers and Shiffrin, 1981). According to relative strength models, the likelihood of retrieving B, given the retrieval cue A, is determined by the absolute strength of the A–B association, relative to the strengths of all associations emanating from A. Such behavior is captured with what is called a *ratio-rule equation*. To illustrate, suppose we wished to compute the probability of recalling the item *rock*, given the cue *dog*, and given that there are N items in total associated to *dog* (including, for example, the item *sky*). A simple ratio-rule equation expressing this recall probability might be written as follows:

p (recall *rock*, given *dog*) = Strength (*dog–rock*)/Strength (*dog–rock*) + Strength (*dog–sky*) . . . Strength (dog–Nth item).

Note that as the associations between the cue and other competing items (e.g., *dog–sky*) become stronger, the probability of recalling the target (*rock*) should decrease, because the denominator of this recall probability equation should increase. Similar effects should occur when additional items become associated to the same cue. Modern theoretical work with architectures using relative strength retrieval assumptions (J. R. Anderson, 1983; Gillund & Shiffrin, 1984; Mensink & Raaijmakers, 1988; Raaijmakers & Shiffrin, 1981; Rundus, 1973) has accommodated a broad array of memory phenomena (see also, Luce, 1959, for a more general treatment of "choice" models). Thus, McGeoch's elementary ideas of competition and strength dependence might provide a general account of interference in the form of an occlusion process (see, e.g., Mensink & Raaijmakers, 1988).

b. Unlearning

Is it possible to forget permanently experiences that have been stored in long-term memory? Consider the following scenario, which does not seem altogether uncommon. An acquaintance visits you and, in the course of reminiscing, describes a conversation that you had at a party several years ago. You may recall, in respectable detail, elements of the party, including your friend's attendance, your various conversations, as well as several amusing events that occurred. However, you may fail to remember having discussed a certain topic specifically with your friend, despite your friend's most confident confirmations, and even when you may clearly recollect both your friend's attendance and having discussed the topic. Subjectively, it seems as though your memory for that past event has become fragmented, impairing your ability to judge how elements of the experience (e.g., your friend and the discussion) go together. One approach to explaining this apparent fragmentation would be to assume that the originally encoded associations between the elements of that event have been disrupted or damaged in permanent fashion by subsequent experience.

Research relevant to the previous intuitions was conducted in the classical interference era in the context of Melton and Irwin's *unlearning hypothesis* regarding retroactive interference (Melton & Irwin, 1940). According to Melton and Irwin, the impaired recall of A–B items typically observed in the retroactive interference paradigm was not merely the result of occlusion of those responses by A–D items; those A–B items were also recalled less well because the associative connections linking As to Bs were weakened during the learning of A–D items. The process of associative unlearning was considered analogous to the extinction of conditioned responses in animal learning—that is, the decrease in the probability of a response, previously conditioned by repeated reinforcements in the presence of some stimulus A, that occurs when that response is subsequently followed by repeated nonreinforcements in the presence of that A stimulus. Melton and Irwin's idea was that whenever the B item intruded inappropriately during practice on the new D target in the A–B, A–D paradigm, the earlier learned A–B association that mediated the offending B intrusion would be unlearned so as to reduce the likelihood of its subsequent intrusion. Because Melton and Irwin combined their unlearning mechanism with the occlusion mechanism postulated by McGeoch, their theory came to be known as the *two-factor theory* of interference.

Melton and Irwin proposed the addition of a separate unlearning mechanism based on perplexing findings observed in a study in which they examined retroactive interference as a function of the degree of interpolated learning. A natural prediction following from McGeoch's occlusion approach is that increases in retroactive interference should be accompanied by increases in overt intrusions of stronger, second-list memory items. Melton and Irwin discovered, however, that as the number of learning trials on a second list was increased to extreme levels (e.g., 10, 20, or 40 learning trials), the frequency of intrusions from the second list decreased, even though retroactive interference continued to increase. Because retroactive interference continued to grow—that is, fewer and fewer of B responses were recalled—even as the frequency of competitor intrusions diminished, Melton and Irwin reasoned that some additional factor, a factor X, must be contributing to retroactive interference. They tentatively identified this factor as unlearning of the A–B association.

Subsequent tests for unlearning yielded supportive results. Because retroactive interference should be produced by both unlearning and occlusion, whereas proactive interference should be produced by only the latter, the theory correctly predicts Melton and von Lackum's (1941) finding that retroactive interference is more pronounced than proactive interference on an immediate test. Support also came from two additional findings. One was that B items suffering from retroactive interference exhibited spontaneous recovery (i.e., the recall of these items got better over time; see Roediger & Guynn, Chapter 7, this volume) as the retention interval between List 2 and the MFR test increased (Underwood, 1948). This finding bolsters the claim that unlearning is analogous to extinction, which also yields spontaneous recovery of the previously conditioned response with increasing delays between extinction and retesting. See, however, Crowder, 1976, for a critique of spontaneous recovery as a signature of unlearning.) The second supportive result was the observation of retroactive interference on the MMFR test (Barnes & Underwood, 1959). This result was initially viewed as especially compelling evidence for unlearning because the MMFR test allowed people generous time to recall both items associated with a stimulus; most researchers presumed such conditions would eliminate the effects of occlusion thought to be at work on tests directing participants to recall only one response (see previous discussion of occlusion).

Investigations of the role of unlearning in retroactive interference that employed recognition memory measures produced less clearly supportive results. According to early theorizing, recognition tests should totally eliminate occlusion because they provide the correct targets, obviating the need for actively retrieving those items from memory (but see our later section on fan effects). As emphasized in our earlier party example, however, unlearning predicts that some retroactive interference ought to be observed even when people are given both the cue and intact target item on a recognition test. Retroactive interference should occur because people's recognition performance depends on their memory of the supposedly unlearned cue–target *associations*. Evidence against the unlearning prediction was obtained by Postman and Stark (1969), who failed to find a statistically significant retroactive interference effect when the nontarget distractor items in the recognition test were other B targets from List 1. Subsequent work demonstrated significant retroactive interference, however, when distractor items on the recognition test were the D targets that had been paired with the tested A cues on List 2 (R. C. Anderson & Watts, 1971). Additional studies favored unlearning, suggesting that intact recognition performance in Postman and Stark's (1969) study may have reflected genuine unlearning of A–B associations masked by participants' use of intact backward associations (B–A, instead of A–B) to match the correct target item with the cue (Greenberg & Wickens, 1972; Merryman, 1971). These studies demonstrated retroactive interference on a recognition test when the backward B–A associations were unlearned as well as the forward A–B associations by including both B–E and A–D learning during List 2. (See, however, Postman & Underwood, 1973, for an interpretation of these findings that discounts unlearning.) Because of the complexity of the arguments highlighted above, it remains an open question whether recognition tests eliminate competition effects, and, as important, whether unlearning actually occurs. (See, however, Wickelgren, 1976, for a systematic listing of studies favoring the unlearning view; see also Loftus, 1979b; Loftus & Loftus, 1980, and our later section on Related Research Areas for discussion of a related view proposed in the context of eyewitness memory research.)

2. Ineffective Retrieval Cues

Suppose that your car breaks down, and you are forced to go to the market in your neighbor's car. If you are like most people, when it comes time to return, you will initially try to remember where you parked your own car. That failing, you will likely realize that you did not, in fact, drive your own car. The failure to recall the target memory in this case arises because your frequent experience of driving your own car to the market led you to use the wrong retrieval cue on this exceptional day. Similar failures might arise for

more remote past events as well. That is, after acquiring considerable experience with a new car, answering queries about prior parking occasions with your old car would be difficult if you failed to recollect that on those occasions you were driving your old car. These are examples of retrieval failures resulting from the use of ineffective retrieval cues—cues that are simply not associated to the target memory. The initial proposal that such failures might account for some instances of interference was called the Variable Stimulus Encoding theory.

Variable Stimulus Encoding Theory

Variable stimulus encoding (VSE) theory was developed by E. Martin (1968, 1971) as an alternative account for the retroactive and proactive interference effects observed in the classical A–B, A–D paradigm. A central assumption of Martin's VSE theory is that two memory responses cannot be simultaneously associated to the same stimulus element (see also Estes, 1955). Because of this restriction, asking people to learn a new set of A–D responses after having learned an initial A–B list forces participants (to the extent that they wish to remember both B and D responses) to encode the repeated A stimuli differently from how they were encoded during A–B learning (hereafter, the List 1 encodings), and to then associate these new encodings of A stimuli (hereafter, the List 2 encodings) with the new D responses. Thus, in a final MMFR test for a given A stimulus, both B and D responses can be given only if the A test stimulus elicits both its List 1 and List 2 encodings, and if both encodings are then used to search memory. If only the List 2 encoding is used, the B response will not be found, resulting in retroactive interference; if only the List 1 encoding is used, the D response will not be found, resulting in proactive interference. Martin assumed that the List 2 encoding of a given A stimulus would be the most accessible encoding immediately after A–D learning, but that the List 1 encoding would become more accessible with increasing delays. Thus, Martin's VSE theory accounts for retroactive interference being greater at short retention intervals, and for proactive interference being greater at long retention intervals.

To illustrate how variable stimulus encoding might produce interference outside the laboratory, consider the following situation. In one of your college classes you meet a casual acquaintance, Sally, with whom you played on a softball team for two years in early grade school, but whom you have not seen since. You once again become casual friends and confine your conversations to college-related activities. Treating Sally as the A stimulus, what you learned about her in grade school can be considered A–B learning, and what you learned in college, A–D learning. Now assume that a mutual college acquaintance asks you for the name of Sally's brother, which you had learned about in grade school—that is, you are asked to recall an A–B association. To retrieve this A–B association, upon being given the A stimulus "Sally," you first use the memory representation that corresponds to your current "image" of her as a young adult and fail to retrieve the earlier learned B response (i.e., "has a brother named Fred"). However, when you then think of Sally as a youngster in her softball uniform, you correctly retrieve her brother's name. In this case, the process of maturation has more or less forced you to have two different encodings of the A stimulus Sally. Martin's VSE theory says that a person in the A–B, A–D paradigm spontaneously shifts from how A was encoded in A–B learning to a different encoding of A during A–D learning. By creating these two different functional encodings of A, the person is able to learn two different responses to the same nominal stimulus.

The early evidence that people can indeed switch stimulus encodings from List 1 to List 2 was somewhat mixed. The theory was supported by results from experiments (Merryman & Merryman, 1971; Richardson & Stanton, 1972; Schneider & Houston, 1968; see Rudy, 1974, for review) that introduced a new stimulus X into each pair in List 2 of the A–B, A–D paradigm (yielding an A–B, AX–D paradigm). For example, *dog–rock* learning would be followed by *dog arm–desk* learning. The results indicated that learners shifted their encoding away from the A stimulus (*dog*) to the X stimulus (*arm*). This resulted in better learning of the X–D (*arm–desk*) association than in a control condition in which the AX–D (*dog arm–desk*) learning was preceded by A–D (*dog–desk*) learning. This shift toward encoding X (*arm*) also resulted in less retroactive interference, that is, better retention of the A–B (*dog–rock*) association than in the standard A–B, A–D paradigm in which no X cue was provided to which encoding could be shifted. Although these results from the A–B, AX–D paradigm indicate that variable stimulus encoding might, in principle, also be operating in the A–B, A–D paradigm, they are not directly relevant to that paradigm. In the A–B, A–D paradigm, no new stimulus component (X) or new encoding of A is explicitly provided by the experimenter during List 2 A–D learning, but must be generated by the learners themselves. To avoid this problem of the A–B, AX–D paradigm, Williams and Underwood (1970) employed trigrams (e.g., three-letter strings, such as XRM) as stimulus terms in a paradigm highly similar to an A–B, A–D paradigm (e.g., XRM–*rock* learning, followed by XRM–*desk* learning). Immediately after XRM–*rock* learning and immediately after XRM–*desk* learning, each individual letter in the XRM stimulus was tested alone. They found that if X was the letter most likely to elicit the B response *rock* after XRM–*rock* learning, it was also the letter most likely to elicit the D response *desk* following XRM–*desk* learning. Hence, contrary to VSE theory, no evidence was obtained to indicate that, in learning a new D response to XRM, the learner shifted to a new letter that had not been associated with the old B response.

Although research directly relevant to VSE theory has largely lain dormant since the mid-1970s, Chandler and Gargano (1995) recently reported a result supporting VSE theory. In their experiment, Chandler and Gargano examined how well people could recall a paired associate such as *child–apple*, given the retrieval cue *child–app_*, after also studying either *child–cookies* or *child–bicycle* in an A–B, A–D paradigm. Compared to a control condition in which only *child–apple* had been studied, studying *child–bicycle* produced the typical retroactive interference effect in the recall of *apple*. However, studying *child–cookies* facilitated recall of *apple*! One interpretation of this finding is that the encoding of *child* during the learning of *child–cookies* was similar to, but not identical with, the encoding of the stimulus term *child* that had occurred in List 1. For example, one aspect of the encoding of *child*, common to both lists, might be "has teeth." This commonality should facilitate *child–apple* to the extent that the persistence of the List 2 encoding encourages an interpretation of the test cue *child* more appropriate for retrieving the List 1 response. During *child–bicycle* learning, on the other hand, *child* might be encoded as *child*–"has legs." This inappropriate encoding of *child* might perseverate into the test, causing people to forget (suffer retroactive interference for) *child–apple*. An additional experiment yielded further support for this VSE interpretation of retroactive interference. In this additional experiment, *child–bicycle* interfered more if it was presented at the time of testing *child–apple* than if it had been studied.

In summary, part of the forgetting that occurs in retroactive and proactive interference paradigms may be mediated by people using the wrong retrieval cue to access memory. Under this analysis, retroactive and proactive interference would occur even though the ability to use the A–B and A–D associations to access responses in memory remains unimpaired (once the appropriate encoding of the A stimulus is selected for searching memory) and even though the availability of the target representation itself remains unaffected by the learning of competing responses. Moreover, the general idea that forgetting may be induced by using different encodings of the cue at study and at test has been applied with good success to many other memory procedures besides the standard interference paradigms (see, e.g., Tulving & Thomson, 1973; Chandler & Fisher, Chapter 14, this volume, and Roediger & Guynn, Chapter 7, this volume).

3. Impaired Target Memories

Sometimes we simply draw a blank when trying to recall something and we have very little confidence that the memory will ultimately be recallable. Such forgetting seems to occur even in the absence of intruding memories, and in the context of a highly appropriate set of retrieval cues. In such cases, the most straightforward explanation of forgetting seems to be that the target memory itself has been lost, disrupted, or otherwise impaired. We consider one such interference mechanism that was postulated in the last stages of the evolution of classical interference theory. This mechanism was called response-set suppression.

Response-Set Suppression

Upon arriving at work, you discover that the management switched your old computer with a new one, which, of course, runs a word processor different from the one you normally use. At first, this shift to a new way of doing things may be difficult. Not only must you learn new commands and procedures for editing your documents, but you must also prevent yourself from trying to do things the old way. This lesson was impressed on one of the present authors, who in hitting the keystroke designating "save" for his old editor lost a document because that keystroke happened to match "quit" in the new software. As time progresses, however, you gradually master the new word processor, and you seem to forget about the old one, including those older features that have no competing function in the new program. It almost seems that one can suppress entire sets of "responses" if their sustained activity substantially interferes with performance on a current task. The proposition that our need to shift "response sets" in this fashion might underlie some forms of interference is known as response-set suppression (Postman et al., 1968).

To see how the idea of response-set suppression can account for retroactive interference, consider the classical retroactive interference procedures from the learner's point of view. The learner must first learn to produce the prescribed set of responses to first-list retrieval cues and is given repeated study–test trials until some level of mastery is achieved at producing those responses. Then, the experimenter gives the learner a second list of items and the task of learning begins all over again. It seems reasonable that people would require a little time to reorient to the new task, particularly if the second list presented the same retrieval cues with different responses. Under these conditions, the tendency inadvertently to provide one of the already "prepared" responses during second-list learning may hinder performance on that task, until the person learns to suppress items from the earlier set and to facilitate second-list responses. Postman et al. (1968) proposed that the facilitation of current responses and suppression of past ones were achieved by a general-purpose "selector" mechanism that operated on the entire class of responses associated with each list. This process was thought to be separate from other mechanisms of associative interference, such as occlusion or unlearning, and to affect the representations of the responses. Thus, response-set suppression attributes some decrement in performance observed in interference studies to changes in the target memories themselves.

The response-set suppression hypothesis was proposed as a means of explaining the conditions under which items suffering retroactive interference would exhibit spontaneous recovery, that is, an absolute increase in their availability with increasing delays (see our earlier section on Unlearning). After its initial demonstration, however, spontaneous recovery proved to be an irregular and variable phenomenon, the conditions of which were unclear (Crowder, 1976). In a detailed analysis of the literature concerning this phenomenon, Postman et al. (1968) proposed that items suffering retroactive interference will exhibit spontaneous recovery when testing procedures favored the maintenance of a "List 2" set, that is, a set that favored the selection of second-list responses and the suppression of the first-list responses. Maintenance of a List 2 set could be enhanced, for example, if testing procedures reminded participants of List 2 responses during tests of List 1 items (e.g., if List 2 responses were presented at test). A List 2 set would be maintained under such conditions, according to Postman et al., because the hypothetical selector mechanism exhibited an inertia that (even without the experimenter's intervention) led people to maintain the most recent response set after it was no longer relevant. According to Postman et al., however, this inertia dissipated over time, allowing earlier responses to recover. Postman et al. showed that when test conditions were manipulated so as to vary the person's tendency to maintain the List 2 response set, the predicted variation in spontaneous recovery could be induced (see also Wheeler, 1995, for a recent and very convincing demonstration of spontaneous recovery).

The response-set suppression hypothesis was consistent with certain recognition memory findings that we discussed earlier. For example, consider Postman and Stark's (1969) demonstration that retroactive interference in the MMFR procedure nearly disappears in a multiple-choice recognition test. Such a finding is readily accommodated by the notion that the List 1 response terms that were being suppressed in the MMFR procedure were now being made available by their presentation in the recognition test. (However, see our section on Unlearning for an alternative interpretation.) Second, response-set suppression can account for the retroactive interference observed in the A–B, C–D procedure. If the previous B response set must be suppressed to facilitate the availability of the D response set during List 2 learning, the perseveration of this List 1 suppression should produce retroactive interference. Furthermore, because no common stimulus terms are shared in the A–B, C–D procedure, neither competition nor unlearning assumptions can account for this finding in a straightforward way (see McGovern, 1964; Mensink & Raaijmakers, 1988, for alternative explanations of this effect).

No theory is perfect, however. Toward the end of the classical interference era, a number of empirical phenomena arose that appeared to be inconsistent with the response-set suppression hypothesis, as initially formulated. For example, when a mixed-list design was used (i.e., both experimental A–D items and control C–E items were randomly intermixed on the same List 2 following A–B learning, rather than having separate A–D and C–E lists studied by different groups of people), retroactive interference was observed only for those A–B pairs for which a corresponding A–D pair was learned on List 2. If retroactive interference is produced solely by a perseverating suppression of the entire List 1 response set, forgetting should have been observed for all of the A–B pairs, not just those for which there was a corresponding A–D competitor. Indeed, Postman and Underwood (1973) conceded that these findings were problematic for the response-set suppression hypothesis and modified their theory to allow for the possibility of stimulus-specific response suppression as well. Crowder (1976), however, justifiably criticized this modification, as it seemed merely to be the unlearning mechanism in disguise, predicting effects virtually identical to those predicted by unlearning. Nonetheless, recent evidence from a new interference procedure, the retrieval practice paradigm, suggests that Postman and Underwood may have been on the right track after all (see later section on retrieval-induced forgetting).

IV. INTERFERENCE IN EPISODIC AND SEMANTIC MEMORY

According to the classical interference perspective, the act of retrieving an item from memory was a matter of eliciting a memory "response" to a "stimulus" cue. This research tradition did not differentiate among qualitatively distinct forms of memory responses. Shortly after the classical interference era, however, Tulving (1972) proposed a distinction between two different varieties of memory, which he argued might be governed by different laws of operation. According to Tulving, answering questions about general knowledge, such as "Who are five people who have received the Nobel Peace Prize?" "How do you pronounce the word *tear*?" and "What are the different meanings of the word *ring*?" taps what he calls *semantic memory*. This form of memory differs from what he referred to as *episodic memory*, which is the type of memory one uses to make temporal/spatial discriminations among episodes that one has experienced, such as the most recent parking of one's car or the presentation of a particular pair of words in a particular list in a psychology laboratory. In this section, we review findings from episodic and semantic memory paradigms that have furthered our understanding of both the mechanisms of interference and the breadth of situations in which interference effects occur.

A. Interference Effects in Episodic Memory

As noted previously, episodic memory refers to one's memory for particular episodes or events that one has experienced at a particular point in the

past. A typical laboratory procedure for examining episodic memory presents people with a list of items to be studied for a later memory test. Items are usually either words or pictures and are often presented only once for a brief period. The later memory test may then be assumed to tap participants' memory for the episode of having seen the items on the study list. (This procedure contrasts with that typically used in the classical interference literature in which people were given an initial training phase composed of repeated study–test trials, designed to ensure participants' ability to elicit the appropriate verbal "responses.") Such episodic memory paradigms have been the primary tools used to explore three more recent interference phenomena that we now describe: part-set cuing inhibition, directed forgetting, and output interference.

1. Part-Set Cuing Inhibition

Most of us have forgotten the name of someone, or something, and have been offered assistance by a well-meaning friend who supplies guesses about the word we are seeking. Unless the friend is lucky and guesses correctly, it often feels as though his or her suggestions make matters worse. Sometimes recall fails until a much later point when, unencumbered by the clutter of incorrect guesses, your mind yields the delinquent name. If you have had this happen, you have experienced, firsthand, the puzzling phenomenon of part-set cuing inhibition.

Part-set cuing inhibition refers to the tendency for target recall to be impaired by the provision of retrieval cues drawn from the same "set" (e.g., category) of items in memory as the target (C. W. Mueller & Watkins, 1977). First, we describe this phenomenon and discuss its core empirical characteristics. (For more comprehensive coverage, see Nickerson, 1984; Roediger & Neely, 1982.) We then present a popular theory that accounts for many of these basic findings and illustrate how this theory, although constructed on the principles of classical interference theory, contributed to a general movement away from the unlearning postulate prevalent in the classical interference era.

a. Basic Findings

The basic phenomenon of part-set cuing inhibition was first nicely illustrated in findings reported by Slamecka (1968). Slamecka had people study lists composed of six words from each of five semantic categories (e.g., trees, birds), which were randomly ordered in the study list. On the final recall test, some people were given some of the members from each category as cues to help them recall the remaining items; others were given no such cues. Of critical concern was people's ability to recall the remaining noncue items in the experimental condition relative to their ability to recall those same items in the condition in which no cues were given. Quite naturally, Slamecka expected that the cues would help recall of the remaining noncue target items. However, when recall was scored for only the noncue target items, people who received the cues performed worse than did people who received no cues! This phenomenon has become known as part-set cuing inhibition (C. W. Mueller & Watkins, 1977) because providing part of the set (which was defined by each semantic category) as cues inhibited performance on the remaining items from that set.

Since Slamecka's initial discovery, several characteristics of part-set cuing inhibition have been consistently observed. In general, as the number of cues given to the subject at recall increases, the ability to recall remaining noncue targets decreases. For example, in an experiment by Roediger (1973), people listened to lists containing varying numbers of exemplars from 16 semantic categories (e.g., Fruits, Trees), blocked by category, and were then given an immediate-recall test in which they were provided with the 16 category names. A critical manipulation concerned whether people received, as additional recall cues, either zero, one, three, or five exemplars from the six-exemplar categories. As the number of these additional exemplar cues increased from zero, one, three, to five cues, the probability of recalling the remaining noncue targets decreased from .66, .63, .59, to .53, respectively. Analogous findings have been observed in a variety of other studies (M. Q. Lewis, 1971; Rundus, 1973; Slamecka, 1968, 1972; Watkins, 1975).

The phenomenon of part-set cuing inhibition is not confined to items selected from the same semantic category. Indeed, in some of Slamecka's (1968) experiments, people heard 30-item lists composed of noncategorized words, varying in their frequency of occurrence in the language and their degree of interrelatedness. At test, they were cued with from 0 to 29 of the list words. A part-set cuing inhibition effect was obtained regardless of word frequency and interrelatedness, presumably because these items merely occurred together in the same experimental context. Part-set cuing inhibition has also been observed by Roediger, Stellon, and Tulving (1977) using unrelated words, and by C. W. Mueller and Watkins (1977) using a variety of "set" definitions, including sets defined by rhyme, by subjective organization, and even by an arbitrary shared cue, as in the paired-associate procedure. Roediger (1978) even demonstrated that cuing people with varying numbers of category names from an earlier-studied categorized word list impaired their free recall of other categories from that list, suggesting that the categories themselves functioned as a set in memory. Thus, a crucial factor underlying part-set cuing inhibition is whether the cues and targets share a common retrieval or "set" cue.

Although part-set cuing makes recall of noncue targets more difficult, it appears to have little or no effect on people's ability to recognize those items. This finding was first demonstrated by Slamecka (1975), who gave people varying numbers of exemplars as cues during forced-choice recognition tests for critical target exemplars. Slamecka found that these cues had no effect on target recognition, irrespective of the number of cues provided. A similar failure to find impaired recognition was observed by Todres and Watkins (1981), although they did find a small effect when nonstudied exemplars served as cues. Because part-set cuing impairs recall more seriously than recognition, it would seem that the deficit reflects a problem in retrieval rather than impairment of the noncue target memories themselves.

b. A Popular Account of Part-Set Cuing Inhibition

On the face of it, part-set cuing inhibition is an extremely puzzling and counterintuitive finding. Why might the presentation of information that is clearly related to the items to be retrieved hurt recall rather than help it? Although a number of factors are likely to contribute to part-set cuing inhibition (see Nickerson, 1984, for a review), the one receiving the greatest attention was proposed in a model by Rundus (1973). In the Rundus model, cuing recall with part of a set impairs performance on the remaining noncue items by inducing retrieval competition between cues and the noncue targets. Thus, Rundus's model applies classical notions of interference (McGeoch, 1942) to explain this intriguing phenomenon.

According to Rundus (1973), people studying a categorized word list encode items in hierarchical fashion with respect to their experimental categories, and those categories with respect to the experimental context. This representation is illustrated in Figure 4, which depicts a potential encoding for the items *orange, banana, grape,* and *apple.* On a recall test, a person first recalls the experimental categories by recalling categories via contextual associations (unless the categories are provided), and then recalls exemplars using each category in turn as a cue. Note here that Rundus assumes that people's memory search is thus guided by a separate *set cue* (or, in his terms, a *control element,* after Estes, 1972)—in this case, a category—in addition to whatever exemplar cues the experimenter may overtly provide. The retrieval process is presumed to be susceptible to strength-dependent competition and is thus modeled in terms of a ratio-rule equation: the probability of recalling an item (e.g., *orange*) to a retrieval cue (e.g., *fruit*) is determined by that item's associative strength to that cue, divided by the strengths of all associations (e.g., *orange, banana, grape,* and *apple*) emanating from that same cue. According to Rundus, presentation of exemplar cues (e.g., *orange, banana*) strengthens those cue items' associations to their shared set cue (i.e., category), reducing the relative strength of noncue targets (e.g., *grape* and *apple*). By this analysis, the strength advantage of cues over noncues causes exemplar cue items to intrude persistently during attempts to retrieve the noncue targets. Impaired recall of noncue targets arises when the number of exemplar cue intrusions exceeds the person's "stopping criterion" for

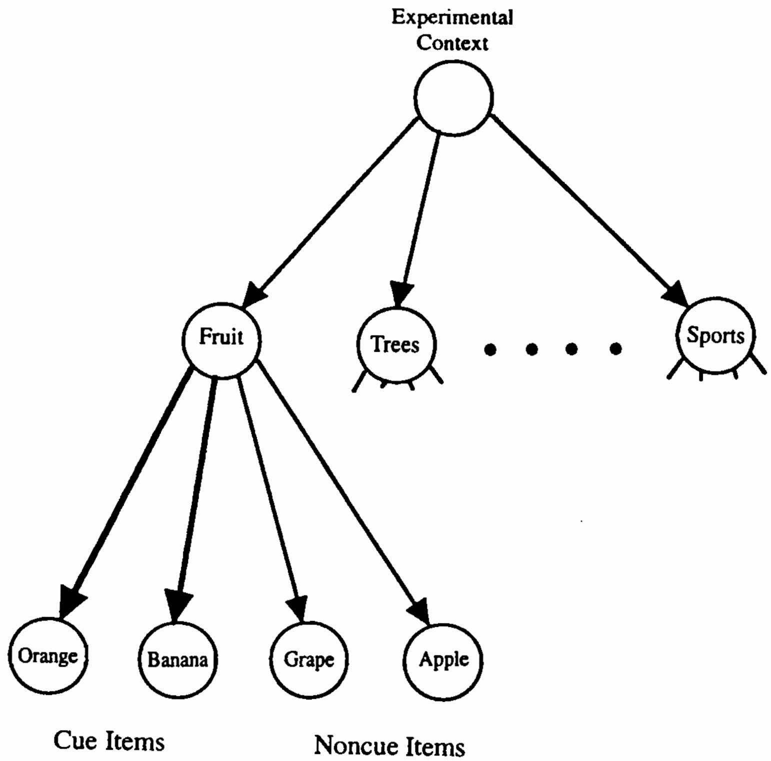

FIGURE 4 A typical representation of a categorized word list that would be assumed by the Rundus model of part-set cuing inhibition. Exemplars of a category are associated to their shared set (category) cue, and categories are associated to a representation of the general experimental context.

recall—essentially the continued intrusion of cue items leads people to give up their memory search. Thus, Rundus's model of part-set cuing inhibition is another example of an occlusion theory (see previous section on occlusion).

The Rundus (1973) model provides a straightforward account of the major empirical findings of part-set cuing inhibition. For example, increasing the number of part-set cues should reduce recall performance for noncue targets because more competitors will have had their relative strengths increased by being presentated as cues. The nature of the set should make no difference as long as the members of the set are associated with a shared set cue (indeed, this seems to be the defining characteristic of "set"). Finally, recognition accuracy should remain unimpaired by part-set cuing because providing the intact target item eliminates the presumed source of impairment: competition for retrieval access. Because the Rundus model provides

such a natural account of these findings in terms of classical principles of competition, it has retained considerable popularity as a means of accounting from this intriguing effect.

The Rundus (1973) model accounts for part-set cuing effects by adapting classical notions of occlusion to this new paradigm. An important contribution of this adaptation was the demonstration that forgetting on unpaced MMFR tests (which is essentially what category-cued recall is) need not reflect unlearning. Recall from our previous discussion of unlearning that the initial observation of retroactive interference on an MMFR test was considered strong evidence of unlearning, because such tests were believed to eliminate occlusion effects (Barnes & Underwood, 1959). The Rundus model illustrated how forgetting on MMFR tests could arise form occlusion alone if one simply assumed that people adopted a "stopping criterion" for memory search, which was exceeded in cases when competitors intruded frequently enough (see also Ceraso & Henderson, 1965, 1966, for earlier challenges of the Barnes & Underwood argument, based on the observation that proactive interference occurs in the MMFR procedure). The ability to account for interference data with occlusion alone, together with the growing body of research showing that interference disappeared (or at least was greatly reduced) on recognition tests, led to a general disenchantment with the unlearning postulate of the two-factor theory of interference. Rather, interference came to be seen as *retrieval inhibition,* that is, a deficit in the ability to retrieve otherwise available memory items. Thus, the discovery of part-set cuing inhibition and its later explanation in terms of occlusion contributed to an important shift that has shaped mathematical and computational thinking about the causes of forgetting. (Later, we discuss a new approach to interference that questions the general viability of occlusion as an account of interference.)

2. Directed Forgetting

So far, we have treated forgetting as a flaw of the cognitive system, an involuntary result of both the structure of our experiences and the properties of memory. In recent decades, however, some work has focused on the benefits of forgetting and on the possibility that memory lapses often arise from voluntary, intentional processes. Consider R. A. Bjork's (1970) example of a short-order cook who during a typical morning breakfast shift must process dozens of highly similar orders. Having completed a particular order such as "Scramble two, crisp bacon, and an English" the cook's performance can suffer only to the extent that the prior orders have not been forgotten. Similarly, we have all experienced times, after completing a memory-demanding activity such as an exam or a well-rehearsed speech, when we wish to "let go" of the information so that our minds may shift to new thoughts and endeavors. When we return to the "dropped" material later, we are often surprised that the material that was so readily accessible only a short time ago now eludes us. These two examples suggest that forgetting may often be an intentional act initiated to reduce the tendency of past experience to impede concentration on some more current activity.

In this section, we highlight basic findings from empirical work on intentional forgetting. Much of this work has been done with what has come to be known as the directed forgetting paradigm (see, e.g., R. A. Bjork, 1972; 1989; Epstein, 1972; Johnson, 1994, for reviews). Other work not discussed here, but relevant to the issue of intentional forgetting has been pursued in research on hypnotic amnesia (Coe, Basden, Basden, & Fikes, 1989; Geiselman, Bjork, & Fishman, 1983; Geiselman, MacKinnon, et al., 1983; Kihlstrom, 1977; 1980; 1983; Kihlstrom & Barnhardt, 1993; Kihlstrom & Evans, 1979); repression (Erdelyi, 1985; 1993; Erdelyi & Goldberg, 1978; Holmes, 1990); thought suppression and mental control (Wegner, 1994; Wegner & Pennebaker, 1993); and on the effects of instructions to disregard certain information in the formation of social judgments (see Johnson, 1994, for a review).

a. Basic Findings

The two most basic, consistently observed findings in the study of directed forgetting are that directing people to forget a set of previously learned items (1) greatly reduces proactive interference from those items on subsequent material, often bringing the level of recall for subsequent information up to the level observed when no prior items are studied, and (2) impedes access to the to-be-forgotten items on a final memory test. These basic findings are nicely illustrated in an early experiment by Reitman, Malin, Bjork, and Higman (1973), who adapted a procedure first introduced by R. A. Bjork (1970). In the Reitman et al. procedure, people received lists of one to eight paired associates, with each associate presented on a computer screen for 2.2 s. During some of these lists, people received a signal to forget all list pairs presented prior to the signal. Following presentation of the list, people received a stimulus member from one of the pairs and were asked to recall the associated response. On most of the forget-cue trials, the test item was drawn from the set of associates appearing after the forget cue. On a smaller number of trials, the test item was drawn from the to-be-forgotten set. Participants were forewarned that such tests of to-be-forgotten items would occur infrequently, and that they would be signaled by an asterisk next to the stimulus member during the tests of those items.

Two features of Reitman et al.'s results are striking. First, recall of items studied prior to the forget cue was impaired by approximately 20%, relative to comparable items on lists in which no forget cue was given. Thus, directing people to forget causes a substantial performance deficit for to-be-forgotten associates. Second, recall of list items presented *after* a forget cue no longer exhibited the typical proactive interference present when people had to remember those same precue items: On lists in which either zero, one, two, or three paired associates were presented prior to a forget signal, people were able to recall 73, 73, 76, and 72% of the postcue associates, respectively—even though performance on those postcue items should have decreased by approximately 8% per precue item without the forget cue (see Reitman et al., Figure 1, for support of this estimate). These data suggest that low recall of to-be-forgotten information was not caused solely by the participants' conforming to the experimenter's forget instructions at the time of retrieval. If such intentional response withholding were the sole cause, it would not have eliminated the typical proactive interference effects of those items when the to-be-remembered items were tested. Thus, when people are directed to forget previous items, they, like Bjork's short-order cook, are able to intentionally forget those items, reaping considerable benefit in the process.

Because people can voluntarily inhibit information they wish to forget so that they may focus on more current tasks, early work on directed forgetting emphasized its potential relationships to repression (Weiner, 1968; Weiner & Reed, 1969). However, most of the early theoretical accounts of directed forgetting attributed the effect to processes that had little to do with inhibition or even forgetting per se (see R. A. Bjork, 1970; 1972; Epstein, 1972, for reviews of these ideas). For instance, evidence suggests that the two main components of the effect—forgetting of to-be-forgotten items, and the consequent reduction in proactive interference on later material—can be explained by factors such as: (1) encoding deficits for "forget" items arising because participants, in anticipation of a forget cue, perform only shallow rehearsal of early items until they know that they will be responsible for remembering them; and (2) differentiation or segregation of to-be-remembered and to-be-forgotten items into discrete sets in memory (see, e.g., R. A. Bjork & Woodward, 1973; Jongeward, Woodward, & Bjork, 1975; D. W. Martin & Kelly, 1974; MacLeod, 1975; Tzeng, Lee, & Wetzel, 1979; Wetzel & Hunt, 1977; Woodward & Bjork, 1971, for evidence bearing on these factors). For example, the case that low recall performance on to-be-forgotten items might simply reflect their poor encoding seems especially plausible in procedures that cue participants to remember or forget *individual words* (as opposed to whole lists)—a conclusion supported by the finding that final recognition performance for individually cued, to-be-forgotten items is consistently inferior to that of to-be-remembered items (Davis & Okada, 1971; MacLeod, 1975; Wetzel & Hunt, 1977; see Basden, Basden, & Gargano, 1993, for a thorough treatment). However, recent work argues that these accounts do not tell the whole story.

b. Evidence for Retrieval Inhibition in Directed Forgetting

Several findings argue that at least part of the impairment observed in studies of directed forgetting derives from a process that impairs access to items successfully encoded into long-term memory. Consider, for example, a classic study by Geiselman, Bjork, and Fishman (1983). Geiselman et al. presented people a list of 48 four-letter nouns, auditorally, with each word preceded by one of two instructions: an instruction either to learn the word in preparation for a final recall test (e.g., learn *hand*) or to judge the word for its pleasantness (e.g., judge *rake*). Midway through the list of 48 words, half of the participants were told that "What you have done thus far has been practice; therefore, you should forget about all of the to-be-learned words that you heard." The remaining participants were also stopped, but were instead told that the first half of the list had been presented and that they should continue to try to remember the to-be-learned words they heard. After the entire list had been presented and a 3-min distractor task had been given, people were given either a recall test for all items—both to-be-learned and to-be-judged words—or a yes–no recognition test containing all learn and judge words together with 48 new distractor items. Geiselman et al. reasoned that if the forget cue caused poorer recall of to-be-forgotten items simply because it induced people to stop rehearsing those items during the second half of the list, then there should be no effect of the forget cue for "judge" words, which the participants were presumably not rehearsing during either half of the list.

People's performance on the "learn" words in Geiselman, Bjork, and Fishman's (1983) experiment showed the typical two-component directed forgetting pattern: participants instructed to forget the initial "learn" words midway through the list recalled fewer first-half learn words (56%) than did people instructed to continue remembering those words (73%), but people instructed to forget recalled more second-half learn words (72%) than did people not allowed to forget the first half (55%). More surprising, however, was the finding that the incidentally encoded "judge" words showed precisely the same two-component pattern, even though participants had no reason to rehearse these items during either half of the list: That is, people directed to forget the initial "learn" words recalled fewer of the judge words from the first list half (30%) than did people instructed to remember first-half learn words (45%), but people instructed to forget recalled more judge words from the second list half (40%) than did those not allowed to forget the first list half (30%). These results indicate that impaired recall performance on to-be-forgotten items in the directed forgetting procedure is more than a failure to rehearse (and thus encode) those items to the same extent as participants instructed to remember those items. This interpretation is supported by the observation that participants in both the forget and remember

conditions, given a final recognition test for all types of words (learn words and judge words from both list halves) recognized all classes of items extremely well, with performance falling between 80 and 85% in all cases. The finding of comparable recognition for forget and remember items argues that to-be-forgotten words are actually encoded into long-term memory, but are rendered inaccessible through some intentional forgetting process, a process that Geiselman et al. referred to as retrieval inhibition. (See Horton & Petruk, 1980, for further evidence, using a levels-of-processing encoding manipulation, that semantically encoded material is subject to directed forgetting; see also Basden et al., 1993, for a replication of the finding of intact recognition memory for to-be-forgotten material when people are cued to forget whole lists of words instead of individual words.)

Additional support for Geiselman et al.'s view comes from an intriguing series of studies examining the circumstances under which this inhibition might be "released." These studies are motivated by the idea that re-exposure of to-be-forgotten material might cause a "rebound" effect in the accessibility of the inhibited information—that is, re-exposure of the information might "release" it from its inhibited state, restoring its accessibility as well as its tendency to interfere with people's ability to recall subsequent material. Findings from a study by E. L. Bjork, Bjork, and Glenberg (1973) support this release-of-inhibition hypothesis. People were presented with word lists (each with 32 items) of three different types: (1) lists with a midlist cue to forget the first list half; (2) lists with a midlist cue to remember the first list half; and (3) lists without a first half, in which the presentation of initial items was replaced by a shape judgment task. After each list, people's memory for the second list half was assessed under one of three conditions: immediate recall, recall delayed by an arithmetic task, or recall delayed by a recognition test for second list half items. E. L. Bjork et al. found performance to be quite similar when recall was tested immediately and when it was delayed by a simple arithmetic task: An instruction to forget the first list half brought people's performance on the second list half (54%) up to the level exhibited by people with no first half (55%), as compared to the clearly inferior performance exhibited by people not allowed to forget those initial items (43%).

More interesting, however, was what E. L. Bjork et al. (1973) found when recall was delayed by an interpolated recognition test for some second list half items. On this interpolated eight-pair recognition test, people were asked to select the item that had appeared on the second list half. On four trials, the distractor item paired with the second list half target was an item from the to-be-forgotten set, instead of a novel, nonexposed item. When recall of the second list half was delayed by this recognition test, people given an instruction to forget the first list half recalled only 35% of postcue second list half items, about as many items as recalled by people required to remember the first list half (33%), and clearly fewer items than were recalled when no initial list half was given (51%). It appears that the mere reexposure of the four first-half distractor items on the recognition test was sufficient to "release" the inhibition of the entire first half list, as measured by the tendency for those items to cause proactive interference during the recall of the second half items. This conclusion was reinforced by the results of a study by E. L. Bjork, Bjork, and White (1984), who replicated the findings of E. L. Bjork et al. (1973), but also found that a recognition test not including first list half distractors did not by itself reinstate proactive interference from the first list half (see R. A. Bjork, 1989, for a discussion of these studies).

Conceptually related work on the "release" of inhibition was reported by Geiselman and Bagheri (1985). They represented both to-be-forgotten and to-be-remembered items for a second study trial (on which all items were then designated as "remember" items). Final recall of to-be-forgotten items benefited more from this representation (39%) than did final recall of the to-be-remembered items (7%). Differential improvement for "forget" items would be expected if those items benefitted both from their repeated encoding and from their "rebound" from their previously inhibited state. When taken together with the E. L. Bjork et al. results, there appears to be intriguing support for the notion of inhibition release.

c. Necessary Conditions for Directed Forgetting

Although the evidence reviewed here suggests that it is possible intentionally to forget previously learned items, there appear to be several limitations on this ability. First, evidence indicates that an instruction to forget is most effective when it follows immediately after the to-be-forgotten items. When a cue to forget is delayed until after additional material has been interpolated, less forgetting is observed for the to-be-forgotten material, and the reduction of proactive interference for later-studied items is smaller or even absent (see, e.g., R. A. Bjork, 1970; Epstein, Massaro, & Wilder, 1972; Epstein & Wilder, 1972; Timmins, 1974, for data bearing on this issue; see also Roediger & Tulving, 1979, for related experiments).[1] A second precondition for effective directed forgetting, suggested by R. A. Bjork (1989), is that new study material must be acquired after the forget instruction is given. Support for this claim comes from a study by Gelfand and

[1] Whether a forget cue can be "aimed" at material farther back in time (e.g., the list that appeared immediately before the most recently studied list) remains to be established. Although some evidence suggests that such delayed forget cues are ineffective, most studies that have examined this issue seem to have confounded this delay with the omission of additional to-be-learned material after the delayed forget cue. Gelfand and Bjork (1985) provided strong evidence that such postcue learning may be a necessary condition for directed forgetting to occur, even under immediate conditions.

Bjork (1985, as reviewed in R. A. Bjork, 1989), in which an initial study list was followed by either (1) an unfilled interval, (2) an unrelated verbal activity, or (3) a second study list. Gelfand and Bjork found that directing people to forget the first list did not impair their final recall of the items in that list when this instruction was followed by an unfilled interval or by unrelated verbal activity (with deficits of 3 and 2%, respectively). When a second study list was given, however, people instructed to forget the first list recalled 17% fewer items from that list on the final recall test, compared to control participants directed to remember those items. These findings suggest that reorientation to new material that substitutes or "replaces" to-be-forgotten items may be necessary intentionally to forget that information. To the extent that the acquisition of new material is essential for the impairment of to-be-forgotten items, the finding of directed forgetting and its interpretation in terms of inhibition begin to resemble the finding of retroactive interference and the hypothesis of response-set suppression (see R. A. Bjork, 1989; Wheeler, 1995, for discussions of this point).

Whatever the relationship between directed forgetting and retroactive interference may be, research on directed forgetting demonstrates that the magnitude of the forgetting observed under conditions of interference depends, in some situations, on people's disposition toward that material. That the magnitude of forgetting depends so strongly on participants' intention to remember (or to forget) suggests that such factors may have played a far greater role than was realized in many classical studies of interference. At the very least, work on directed forgetting provides a precedent illustrating that forgetting may sometimes be intentional and controllable, and that such forgetting may have significant advantages for the current focus of cognition. We return to a related perspective in our later discussion of retrieval-induced forgetting.

3. Output Interference

If you are fond of constructing lists—such as a list of things to do, items to buy, or people to invite to a party—you have probably experienced the sensation that generating new items for your list gets more difficult as you proceed. The most natural interpretation of this sensation is that you are, in fact, running out of things to list, and that if you leave something off, it must not be particularly important. Although there is some truth to these intuitions, we often do omit important things. Such omissions are likely to be a product of what is known as output interference.

Output interference refers to the gradual decline in memory performance as a function of an item's position in a testing sequence. For example, a person's ability to recall the word *sky* in response to the cue *dog* will decrease if that item is tested later rather than sooner in a testing sequence. This decline in recall performance with testing position was first observed in a study using paired associates (Tulving & Arbuckle, 1963) and was originally attributed to the loss of information from short-term memory over the interval of testing. However, a number of findings show that output interference occurs even when the contribution of short-term memory to performance is eliminated. For example, giving people an unrelated task to occupy their short-term memory in the interval between the initial study and the final recall test does not affect the degree of output interference (A. D. Smith, 1971). Further, output interference does not depend on the position of a category in the study list (or even on the position of an item within a category set; see Roediger & Schmidt, 1980). If the loss of items from short-term memory were responsible for output interference, the drop-off in recall should have been worse when analyses focused on the recall of later (more recent) learned categories rather than earlier-learned categories (if one assumes that short-term memory contributes more to recall performance for later categories). Thus, the crucial variable modulating output interference appears to be the amount of prior retrieval, not the passage of time.

An intriguing characteristic of output interference is that it appears to violate the widely held idea, recurring throughout this chapter, that interference is initiated by competition for a shared retrieval cue. In the output interference procedure, recall of a target item is impaired by the previous retrieval of other items whether or not the target shares cues with those retrieved items. Consider a study by A. D. Smith (1971), in which people studied seven items from seven unrelated semantic categories. On a final recall test in which the people were cued with each category name in turn, the average number of items recalled per category dropped systematically from approximately 70% for the first category tested to 45% for the seventh category tested. Roediger and Schmidt (1980) obtained analogous findings with paired associates. After studying 20 pairs, people were given the stimulus term of each associate as a cue for the recall of its target. Across the five sequential test blocks (each block containing four test cues), there was a systematic decline in the probability of correct recall (.85, .83, .80, .76, and .73, respectively.) Thus, the decline in recall caused by prior output is nonspecific in that it extends across "sets" and does not depend on set type (to use the language of part-set cuing). This cross-cue impairment resists straightforward interpretation in terms of competition for a shared cue, although one might appeal to a more generalized competition for an experimental context cue (see, however, M. C. Anderson & Spellman, 1995, reviewed later, for evidence against this interpretation).

Although most studies of output interference have employed recall tests, the phenomenon has also been observed in recognition. For example, A. D. Smith (1971) had people study a list containing seven categories and assessed recognition memory for targets from those categories, as a function

of the category's position on the final recognition test. The test provided each category name together with seven studied target exemplars and seven nonstudied distractors, requiring people to decide whether each item had appeared on the initial study list. Correct recognition declined across the first three test positions, leveling off for the remaining positions. Similar findings have been observed in other studies (see, e.g., Ratcliff, Clark, & Shiffrin, 1990; Ratcliff & Murdock, 1976). Thus, output interference does not appear to be specific to the task of cued or free recall.

Output interference has also been demonstrated in a procedure that measured participants' recognition time in addition to recognition accuracy. In an experiment by Neely, Schmidt, and Roediger (1983), people studied a categorized list with five exemplars per category (targets) and then made speeded recognition judgments to targets and distractor exemplars taken from those categories. Neely et al. varied whether the critical target item was preceded by (1) a test item from the same category or from a different category, and (2) two or six test items from the same category as the item immediately preceding the target. In addition, Neely et al. equated the nonspecific cross-category output interference effects discussed previously by measuring performance on critical test items occupying the same overall positions in the test list. The foregoing manipulations had only small effects on recognition accuracy, but clearly influenced how fast people made recognition judgments: People were faster when the preceding item was from the same category (as compared to an unrelated category), but were slower when the critical test item was preceded by six rather than two same-category test items. This result illustrates that, holding generalized output interference effects constant, one can observe category-specific output interference on recognition speed.

Two general implications of output interference are important. First, output interference clearly demonstrates a case of interference among items that do not share retrieval cues. Such interference appears inconsistent with straightforward accounts of the impairment in terms of competition among items that share the same retrieval cue—a point to which we return in our later discussion of cue-independent forgetting. Second, output interference shows that the retrieval situation itself might be a source of forgetting. To understand the importance of this observation, one need only consider the ubiquity of this basic cognitive process in our daily cognitive experience. That is, any cognitive act that makes reference to representations stored in memory (which is likely to be all processes) employs retrieval. If retrieval is a source of interference, then accessing what we already know might contribute to forgetting, independent of the encoding of new experience. This implication was first emphasized by Roediger (1974), although its ramifications were not fully appreciated by others. However, this observation forms the starting point for the most recent perspectives on the nature of interference, described in the final section of this chapter.

B. Interference Effects in Semantic Memory

The distinction between episodic and semantic memory led to a large body of research exploring the characteristics of people's memory for very well established, general knowledge. A typical procedure for examining such general knowledge might measure how fast people can make various judgments about an item in semantic memory, such as judgments about the truth of a fact (e.g., deciding whether birds have wings), or about the status of an item as a word (e.g., deciding whether *dog* is a word). Often, the aim of this research is to examine *how* (as opposed to *whether*) retrieval is performed, with considerable analysis given to the on-line dynamics of retrieval such as the rate, extent, and longevity of spreading activation within the semantic network. In this section, we discuss research demonstrating that interference effects occur even in tasks such as these that tap very well learned general knowledge. We discuss two areas of research in semantic memory that illustrate the breadth of situations in which such effects occur: Interference in retrieving facts and interference in our ability to understand a word's meaning.

1. Fact Retrieval, Fan Effects, and the Paradox of Interference

As E. E. Smith, Adams, and Schorr (1978) pointed out, if semantic memory for facts were as susceptible to interference as episodic memory, one would be confronted with what they called the paradox of interference. Specifically, as an expert learned more and more facts about a given topic area, he or she should develop more and more difficulty in remembering any one of them. Although this does not seem to happen in real life, early studies on the speed of verifying "facts" learned in the laboratory showed interference effects (J. R. Anderson, 1974). After students intensively studied many "facts" about various people being in various locations such as "A hippie is in the park," "A hippie is in the church," and "A lawyer is in the school," they were asked to verify "A hippie is in the park" was true or "A lawyer is in the church" was false. An interference effect was observed in that the more different facts that were learned about a person or location, the longer it took for the students to verify a statement about that person or location.

To account for this interference, J. R. Anderson (1976) assumed that the facts the students learned were stored in memory as a network of associations, such as the example depicted in Figure 5. When a test sentence such as "A hippie is in the park" was presented, the memory nodes corresponding to *hippie* and *park* would be activated, and this activation would spread down the links emanating from these nodes. If the activation spreading from one node intersected with activation spreading from the other node down the associative link connecting them, a "true" response would be made. Under the assumption that a node's capacity for sending out a wave

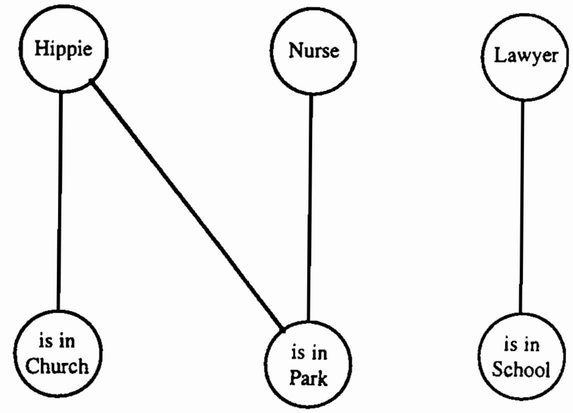

FIGURE 5 A simplified network representation depicting the facts "The hippie is in the church," "The hippie is in the park," "The nurse is in the park," and "The lawyer is in the school." The concept "hippie" is associated to (i.e., fans out to) two locations, and "park" fans to two people, decreasing the effectiveness of both of these concepts as cues, relative to either "lawyer" or "school" or "nurse," each of which is associated to only one thing. These representations are simplified from those proposed by J. R. Anderson (1974).

of activation is resource limited, the speed at which activation spreads down an associative link emanating from that node would be slower the greater the number of other associative links fanning out of that same node. The more different facts the person learns about someone, such as the hippie, the more links there are fanning out from the node for hippie, and the longer verification times become, a phenomenon known as the *fan effect*. Thus, Anderson's model of the fan effect assumes that associations emanating from concepts in semantic memory compete for activational resources, consistent with the broader competition assumption sketched earlier.

Of course, J. R. Anderson's (1974) "fact"-retrieval experiment might not really have tested semantic memory rather than episodic memory. That is, students might have performed this memory task not by "verifying facts" but rather by determining whether the test sentence had been studied at a particular time and place, namely, in the memory experiment. To examine this issue, C. H. Lewis and Anderson (1976) had students study both true facts ("Teddy Kennedy is a liberal senator") and up to four fantasy facts (e.g., "Teddy Kennedy wrote Tom Sawyer") about well-known people. Even when the students knew that they were being tested only on true facts and hence did not need to refer to what they had learned in the experiment to perform well, a fan effect occurred as the number of fantasy facts learned about the famous person increased. If the learning of only four additional facts about someone about whom many facts are already known produces memory interference, then the number of facts learned by an expert should produce massive interference. Yet experts retrieve facts quickly. How can this paradox be resolved?

a. Resolving the Paradox of Interference

McCloskey and Bigler (1980) and Reder and Anderson (1980) were able to provide a partial resolution of the paradox of interference. Their experiments showed that by focusing memory search on only information specifically relevant to a memory query, one can restrict the source of interference to only the small number of other directly relevant facts and greatly reduce or eliminate interference from all of the other facts stored about that general topic. To borrow an example from McCloskey and Bigler (1980), an expert on Richard Nixon might have stored different subcategories about Richard Nixon, such as his foreign policy views, his family life, and his Watergate actions. If asked a question about Nixon's wife, this expert would not search through all of the facts she or he knows about Nixon, but rather would search only through those facts relevant to his family, with only the facts stored under that subcategory producing memory interference. McCloskey and Bigler (1980) and Reder and Anderson (1980) independently proposed and tested this idea in a series of clever experiments and obtained similar results. (See also Bower, Thompson-Schill, & Tulving, 1994, for a similar type of experiment using an MMFR test in the classical A–B, A–D paired-associate paradigm.)

To create different "subcategories" on which memory search could be focused, Reder and Anderson (1980) used narrative materials in which named people (e.g., Alan) performed a series of actions relevant to different scenarios, such as taking a train trip or going skiing. Figure 6 displays how these materials were presumed to have been stored. People who studied the materials in (C–F) learned about Alan participating in both scenarios, whereas people who studied the materials in (A) and (B) learned only about Alan taking a train trip. The foils (test items that were false) were always related to the true test items in that they were always related to the scenario(s) in which Alan had participated. For true facts, such as "Alan checked the weekend Amtrak schedule," verification times revealed a fan effect as the number of facts learned about Alan's train trip increased (A vs. B; C vs. D; E vs. F), but not as the number of facts learned about Alan skiing increased (C vs. E; D vs. F). However, verification times for statements

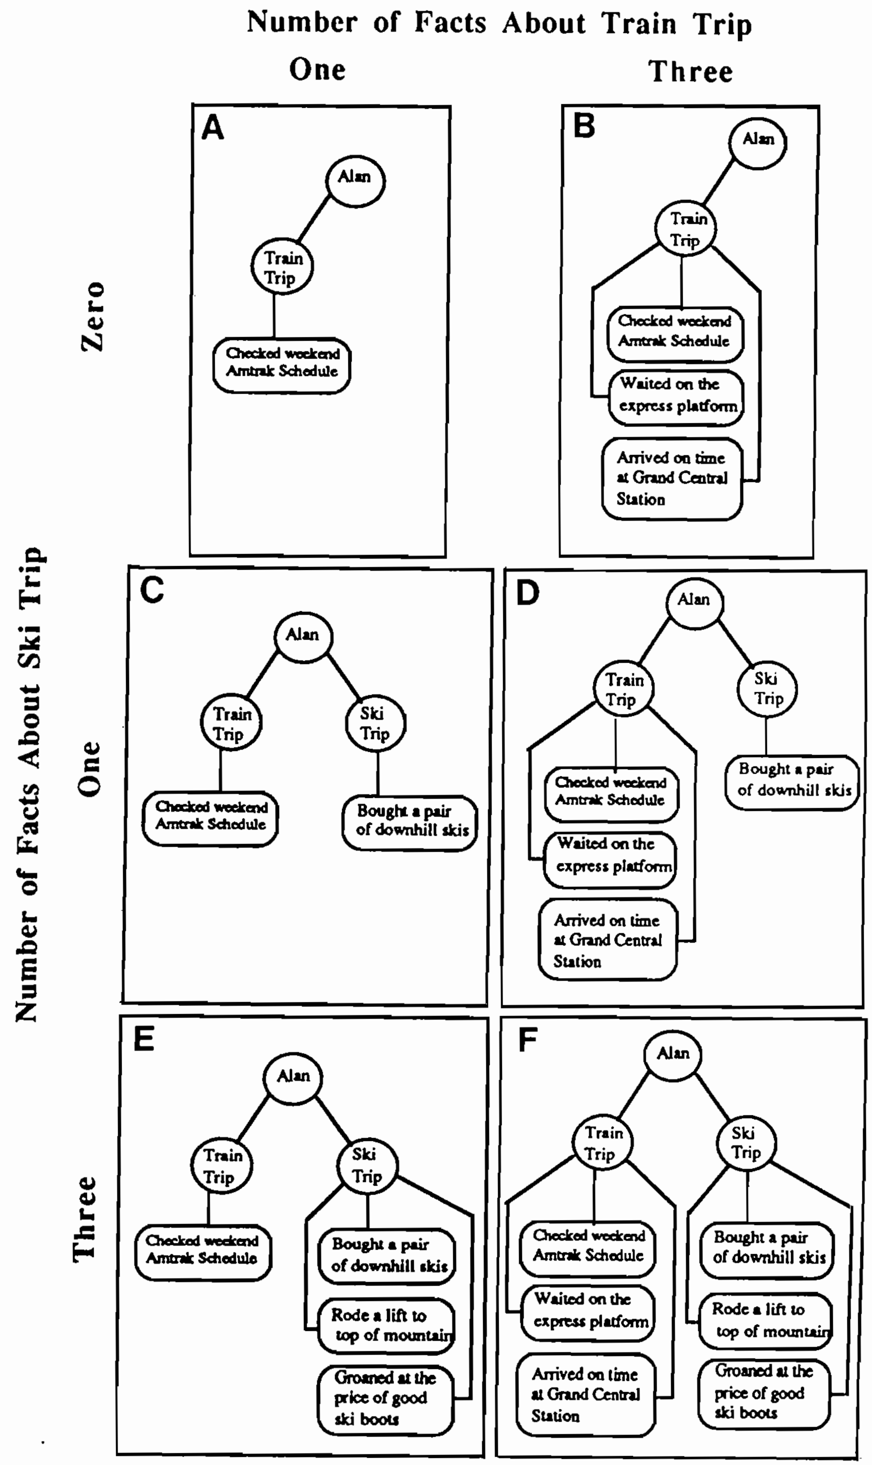

FIGURE 6 Simplified network representations for the materials in each of the conditions of the Reder and Anderson (1980) fan-effect experiment. Subjects studied either zero, one, or three facts about a fictional person named Alan going on a ski trip (spanning top to bottom), and either one or three facts about Alan going on a train trip (spanning left to right). Subjects are assumed to represent facts about each of these separate trips in distinct "subcategories" (depicted in the figure as nodes lying intermediate between Alan and scenario facts) that reduce memory search complexity.

about Alan's train trip were slower if Alan had also participated in the skiing scenario (C–F) than if he had not (A and B). These results imply that people first select the relevant subcategory, and the time it takes to do this is affected by the number of irrelevant subcategories but not the number of facts learned within those irrelevant subcategories. Once the relevant subcategory is selected, only the number of facts within that category influences search time. McCloskey and Bigler (1980) obtained similar fan effects when search could be restricted on the basis of whether the grammatical object of a to-be-verified fact was an animal or a country, and all of the foils were related to the true test items, that is, were about animals or countries.

It is important to note that the fan effects produced by items within the relevant scenario were obtained when all of the foils were related to the true test items, as was described earlier. However, Reder and Anderson (1980) showed that even a fan effect from items within the relevant scenario can be reduced when the foils are unrelated to the true test items. Reder and Anderson argued that with unrelated foils, a person can accurately verify a test item merely on the basis of its plausibility. To understand this, consider the conditions represented in (C–F) of Figure 6, in which all of the true items are about Alan's taking a train ride and going skiing (which he did) and all of the foils are unrelated to these scenarios (e.g., are about Alan going to the circus, which he did not do). To respond accurately to a "true" item, such as "Alan arrived on time at Grand Central station" or "Alan groaned at the price of good ski boots," or to a false item, such as "Alan liked the trapeze artists the best" one would need to decide only if each statement is plausibly true about Alan (i.e., is related to taking a train ride or going skiing) and would not need to look up that particular fact. Under such circumstances, one must still determine in which general scenarios Alan had participated, such that two scenarios (C–F) would lead to slower verification times than only one scenario (A and B). Indeed, this effect was still observed.

However, once people determined that one of the studied scenarios was or was not the one being tested, they could immediately respond "true" or "false," respectively, without searching for the specific fact within that scenario. Thus, the fan within that relevant scenario would no longer slow down verification times, thereby accounting for the greatly reduced fan effect within the relevant scenario. Indeed, under some circumstances in which plausibility judgments can be used, and especially when the retention interval is long, verification times for "true" statements are faster, not slower, the more facts that have been learned about the tested subcategory or scenario (Reder & Ross, 1983; Reder & Wible, 1984). This speed up (negative fan effect) occurs because the more facts that have been learned about the relevant subcategory or scenario, the greater is the total activation level of the category node that is being retrieved in making the plausibility judgment, thereby speeding its access. (See Myers, O'Brien, Balota, & Toyofuku, 1984; Reder, 1982, for the roles that the use of integrated causally linked facts and plausibility, respectively, play in producing negative fan effects. See also Radvansky & Zacks, 1991, who have shown that fan effects from relevant facts may be eliminated even when a plausibility judgment cannot produce accurate performance. This result can occur when learners can create a representation of many objects residing in the same location, what Johnson-Laird, 1983, has called a mental model. Finally, see Conway & Engle, 1994, for evidence that fan effects in long-term memory do not vary as a function of an individual's short-term working-memory capacity; see also Nairne, Chapter 4, this volume.)

b. Implications of Fan-Effect Research

What are the implications of fan effects for our present discussion of interference? First, they illustrate clearly that interference effects occur in both semantic and episodic memory. Thus, although Tulving's (1972) speculations about the differences between semantic and episodic memory may be true for other aspects of these systems, the basic properties of interference appear similar for these two forms of memory. However, interference effects that occur in tests of "semantic" memory can be greatly reduced in at least two ways. Interference may be reduced when memory is compartmentalized into subcategories that allow for a focused search that restricts the source of interference to only those items within that subcategory (but see Whitlow, 1984, for potential limitations of this analysis); or interference may be reduced when correct responses can be made merely on the basis of the test item's plausibility. In fact, when plausibility judgments suffice for responding correctly, response times can be faster the more different facts have been learned about the topic—a negative fan effect. Such findings help explain how experts avoid interference and retrieve information so quickly. A second implication is that fan effects in semantic memory demonstrate that interference can occur in recognition tests. Recall that according to classical interference theory, recognition tests should eliminate response competition effects. By that analysis, either unlearning or response suppression would be the source of these interference effects. However, this implication does not hold in the context of more recent theories (e.g., J. R. Anderson, 1976) in which multiple associative links emanating from the same memory node compete for the flow of a resource-limited spread of activation from that node.

C. Interference in Word Meaning Retrieval

The research on fan effects discussed in the previous section illustrates how considerations of interference enter into the ostensibly simple task of judging the truth of a fact. Interference effects can occur on an even more basic level than this, however. For instance, in the English language, many sources of interference can be found even in apparently effortless, highly practiced tasks like retrieving information about a word's pronunciation, spelling, or its meaning. Retrieving the pronunciation of a word is potentially subject to interference when, for example, whole words like *tear* are associated with two distinct pronunciations (one associated with an action that can rend cloth and the other associated with crying). Similarly, interference can arise in retrieving a word's spelling when a sound like /dō/ is associated with more than one spelling, that is, *doe* and *dough*. Because two different "responses" must be associated with the same linguistic stimulus in each of these cases, the learning that occurs for such items represents the perfect semantic memory analogue of a randomly intermixed A–B, A–D paradigm.

An important example of interference in language processing occurs in the retrieval of word meanings, particularly in the case of homonym meaning retrieval. Homonyms are words with at least two distinct meanings associated with the same spelling and pronunciation (e.g., the word *ring*, with one of its meanings being related to jewelry and the other to a bell's sound). In both speeded lexical (word vs nonword) decision and pronunciation tasks, response times to homonyms are typically faster than to words having only one meaning (Balota, Ferraro, & Connor, 1991; Joordens & Besner, 1994). However, because one need not discriminate between the homonym's different meanings to perform these two tasks, this facilitative effect of semantic ambiguity may be viewed as analogous to the negative fan effect observed for plausability judgments. More relevant to potential interference effects is how a person retrieves the appropriate meaning from a homonym's multiple semantic interpretations. To determine which of the homonym's two meanings is activated at various times after meaning retrieval has begun, researchers have used the *semantic priming paradigm*. In this paradigm (Neely, 1976), people are asked to respond as quickly as possible to a target word (e.g., *cat* by either pronouncing it or making a lexical decision to it. Immediately preceding this target, a word (called the prime) is presented. Not surprisingly, people respond more quickly to the target when it follows a semantically related prime (e.g., *dog*) compared to an unrelated prime (e.g., *wall*).

The semantic priming paradigm has been used to examine the time course of activation for dominant (frequent) and subordinate (less frequent) interpretations of homonyms. For example, Burgess and Simpson (1988) presented homonyms (e.g., *ring*) as visual primes. Either 35 or 750 ms after a prime, a target was presented to the left or to the right of a fixation point to manipulate which cerebral hemisphere of the brain would process the target first (see Springer & Deutsch, 1993, for review). For present purposes, the most interesting result occurred for targets that were first processed by the left hemisphere. For targets related to the homonym's dominant meaning (e.g., *diamond*), priming occured at both delays. For targets related to the nondominant meaning (e.g., *bell*), however, priming was restricted to the 35-ms delay, with the 750-ms delay exhibiting an inhibitory priming effect: Responses were now slower to a target related to the homonymic prime's nondominant meaning than to a target preceded by a totally unrelated prime. These data suggest that when a homonym is recognized, the left hemisphere selectively focuses attention on its dominant meaning and actively suppresses its nondominant interpretation. (See Marcel, 1980; Simpson & Kang, 1994, for evidence of similar suppression effects that occur when the homonymic prime and target word are presented at visual fixation.)

These studies of the dynamics of homonym meaning retrieval illustrate two important points about memory interference. First, they show that interference can occur even for the highly overlearned and seemingly effortless memory retrieval involved in accessing an individual word's meaning. These findings make it abundantly clear how ubiquitous interference phenomena are in cognition, influencing performance across a variety of cognitive tasks. Second, these data on the time course of meaning retrieval support the proposal that suppression mechanisms contribute to changes in the accessibility of knowledge in semantic memory. If such retrieval processes influence the accessibility of items in semantic memory, the possibility arises that similar mechanisms may influence the accessibility of items more generally. This issue is the concern of our next section.

D. Retrieval-Induced Forgetting: A New Perspective on Interference

Studies of episodic and semantic memory have yielded a number of important insights and findings that pertain to the mechanisms of interference. At least two of these insights have contributed to a recent perspective on the causes of forgetting. First, studies of output interference and part-set cuing have highlighted how the act of recall itself might be a source of forgetting in episodic memory (see Roediger, 1974, 1978, for clear proposals of this view). Second, the more detailed analysis of the retrieval process that accompanied research on semantic memory has emphasized how retrieval was not simply a matter of "responding" to a stimulus, but was a complex process the dynamics of which could be examined empirically. These insights form the foundation of a new perspective on interference in which fluctuations in the accessibility of information in both semantic and episodic memory derive from suppression mechanisms that are tied to the retrieval process itself. To take the example offered in the introduction to the present chapter, you do not forget where you parked yesterday because storing subsequent parking memories alters the representation of yesterday's episode; rather, such forgetting stems from your suppressing yesterday's parking episode while retrieving where you parked today. Thus, the impairment associated with the learning of interfering materials is seen as a problem of *retrieval-induced forgetting*.

What are the dynamics of the retrieval process that might be responsible for the impaired recall observed in studies of interference? Consider the following analysis of the functional circumstances often faced during retrieval tasks. Retrieval ordinarily begins with a cue that is necessarily incomplete as a specification of the target memory. For instance, remembering that we are supposed to buy some fruit is a helpful start to an outing at the market, but *fruit* is overly general as a cue if we wish to remember to buy oranges. Such general cues will necessarily be consistent with many potential targets in memory (e.g., *lemon, banana*) that, we might assume, also become active in response to that cue. When activation spreads broadly in this manner, retrieval competition will ensue and access to the target item will be momentarily impeded. If the resolution of such competition in favor of the target item were achieved by a suppression process that focused activation to target items, the consequences of that inhibition should be observable as a decrement in the performance on the inhibited item on subsequent tasks. That is, accessing target items may entail suppression of competitors that can be seen as retrieval-induced forgetting of those competitors on later recall attempts. In this section, we review recent work on retrieval-induced forgetting in both semantic and episodic memory paradigms that supports the operation of a special retrieval-based suppression process causing interference effects in subsequent retrieval.

E. Retrieval-Induced Forgetting in Semantic Memory

Have you ever tried to retrieve a word, fact, or name from semantic memory and felt on the verge of being able to do so, only to fail because some other related item comes to mind and seems to block your retrieval of the item you are trying to recall? As noted earlier when we introduced occlusion, R. Brown and McNeill (1966) have called this experience the tip-of-the-tongue state, a state that apparently occurs rather frequently (at least as reported in Reason & Lucas, 1984, who had people keep track of and classify the tip-of-the-tongue states they had in their everyday lives; see also A. S. Brown, 1991). We now discuss experimental data indicating that such retrieval-induced forgetting in semantic memory may be produced not by a retrieval block induced by the conscious retrieval of related interlopers, as the subjective experience associated with the tip-of-the-tongue state suggests, but rather by an active suppression mechanism that operates when retrieval is difficult.

1. Evidence for Impairment Specific to Retrieval

Perhaps the earliest demonstration of retrieval-induced forgetting in semantic memory may be found in a study of part-list cuing inhibition by J. Brown (1968). Brown had one group of people study a list of 25 American states for a period of five min, while a control group did light reading. On an immediate recall test, he instructed both groups to list all 50 of the American states. Of critical concern was participants' ability to recall the remaining noncue states from semantic memory in these two conditions. The people who received the cues recalled the remaining noncue states more poorly than did those who received no cues, even though the recall period for both groups extended for a full 10 min. One might argue that the impaired recall of noncue items by the cued group reflects retrieval-induced forgetting because the cued group recalled more of the cue states earlier in the output sequence than did the noncued group. Evidence supporting this speculative interpretation of Brown's data comes from Karchmer and Winograd's (1971) demonstration that the cuing inhibition effect was accentuated when people were explicitly instructed to retrieve the states that had served as cues first. Thus, these findings may be taken as early evidence that prior retrieval of items from semantic memory impairs retrieval of related items.

A more systematic examination of retrieval-induced inhibition in semantic memory was undertaken by A. S. Brown (1981), who attempted to control the nature of people's prior retrievals. In Brown's experiment, people were presented with a category name and a letter (e.g., *fruit*–g) and asked to report an exemplar of that category beginning with that cued letter (e.g., *grape*). Each category name was followed by five exemplar trials, one at a time, at about 5-s intervals, with each exemplar cued by its own first letter. The time to retrieve an exemplar increased from the first to the fifth retrieval within that category, suggesting that prior retrievals from a semantic category induce "forgetting" (as represented by slower retrieval times) of a subsequently retrieved target item from that same category. However, it is unclear from this result whether it was the attempt actively to retrieve the prior "cuing" items or their mere (albeit self-) presentation that was responsible for slowed target retrieval.

To determine whether or not the prior cued retrievals were responsible for the impairment in A. S. Brown's (1981) experiment, Blaxton and Neely (1983) directly compared the effects of prior retrieval and prior presentation of competitors on target retrieval speed. In Blaxton and Neely's study, people either actively generated or read aloud a category exemplar on each of the "cuing" trials preceding the critical target trial. On the target trials of interest here, people were to generate a category exemplar in response to a letter cue, as in the Brown study. As the number of cue exemplars preceding the target trial increased from one to four, retrieval times for the target increased, but only when cues were actively retrieved and were from the same semantic category as the target item. Thus, retrieval of the cue items appears to be directly responsible for the impairment of target items in this paradigm because the mere presentation and reading of cues did not hinder the generation of the target item.[2]

Although Blaxton and Neely's (1983) study clearly illustrates a case of retrieval-induced forgetting in semantic memory retrieval, one might still argue that these data do not necessitate the postulation of retrieval-based suppression processes of the sort relevant to the present discussion. Indeed, Blaxton and Neely suggested two explanations of their findings. First, the slowed recall of target exemplars might arise from the blocking of those critical items by the highly available exemplar competitors that had just been retrieved on prior cue trials. For example, given the target trial *fruit*–g, people might retrieve *apple* and *banana* before retrieving the target *grape*. To explain the absence of the hypothesized blocking when primes were merely presented, one might assume that actively generating cue items makes those items more strongly competitive than merely reading them aloud. Thus, impairment might not have been observed in the reading condition because competition from cue items was not strong enough. Alternatively, the slowed recall of target exemplars might reflect the suppression of those target items that occurred during the previous retrievals of cue exemplars. Under this account, *grape* might have covertly intruded during the previous *fruit* cuing trials, rendering *grape* vulnerable to retrieval-based suppression processes. The slowdown in target retrieval as the number of cue trials increased follows naturally if one simply assumes that *grape* would have been suppressed more often with four previous *fruit* cuing trials than with one. However, because both of these explanations can account for the Blaxton and Neely findings, these data only support, but do not demand, the postulation of an active suppression process.

[2] It is important to mention that the interference effect that Blaxton and Neely (1983) reported depended on there being multiple retrievals from the same semantic category. When only one item had been actively generated from the same semantic category as the target, target retrieval was facilitated, relative to when there was active generation of one item from an unrelated category. This facilitation effect conceptually replicated earlier results by Loftus and Loftus (1974). Moreover, a similar facilitation effect occurs in a definition answering paradigm, in which people are asked to retrieve a relatively rare word (e.g., banshee) from its definition (e.g., female spirit whose wail portends death), when a single semantically related cue (ghoul) is read before or after the definition. (See Roediger, Neely, & Blaxton, 1983, and Meyer & Bock, 1992, the latter of whom also showed that the related cue induced more tip-of-the-tongue states.) However, this latter facilitation effect occurs only if the cue that is read is never the correct answer to the definition. When the cue sometimes is the correct answer, the semantically related cue slows target retrieval (A. S. Brown, 1979; Roediger, Neely, & Blaxton, 1983).

2. Evidence for Impairment Caused by Suppression

Recent work by Carr, Dagenbach, and colleagues (Carr & Dagenbach, 1990; Dagenbach & Carr, 1994a; Dagenbach, Carr, & Barnhardt, 1990; Dagenbach, Carr, & Wilhelmson, 1989) argues more clearly in favor of a suppression mechanism mediating retrieval-induced forgetting in semantic memory. In the Dagenbach et al. (1990) study, people learned the meanings of obscure vocabulary items such as accipiter (i.e., hawk) and subsequently participated in a lexical decision task in which these newly learned words served as primes. For the lexical decision task, people were asked to try to retrieve the meaning of the prime word during its 2-s presentation, and to use this meaning to predict what the lexical decision target would be. After the lexical decision test, Dagenbach et al. both tested people's ability to recall the meaning of the primes and administered a recognition accuracy test. They then limited their analysis of lexical decision performance to those trials containing primes for which people could correctly recognize but could not recall the meanings. The presumption was that if the prime's meaning could not be recalled in the unpaced recall test, retrieval failure would have occurred during the 2-s interval allowed during the lexical decision task.

Dagenbach et al.'s (1990) analysis of their lexical decision data clearly supports the notion that a retrieval-based suppression process may be invoked when retrieval is difficult. When people failed to retrieve the meaning of a prime word, such as *accipiter,* lexical decisions for a related target, such as *eagle,* were markedly slower, relative to decisions made for an unrelated target such as *clam* under those same conditions. When people could successfully recall the prime word's meaning, however, lexical decisions on related targets were facilitated. Thus, when people struggled (and failed) to retrieve the meaning of an item from memory, information related to the sought-after item suffered an increase in reaction time, consistent with the notion that those related items were suppressed when they interfered with the difficult retrieval. Dagenbach et al. (1989) obtained a similar result when well-known words (e.g., *cat*) served as primes, provided that meaning retrieval was impeded by making the prime hard to see through brief presentations and masking. As in the Blaxton and Neely (1983) study, when conditions encouraged people actively to retrieve the prime's meaning, related primes produced inhibition; when conditions encouraged passive processing of the prime, related primes yielded facilitation. Neither of the present findings is easily accommodated by blocking mechanisms, because target impairment occurred even when the putative blocking items were never successfully retrieved.

To account for why retrieval failure might impair the subsequent retrieval of related information, Carr and Dagenbach (1990) postulated that when one fails to retrieve a word's meaning under arduous conditions (e.g., when the meaning is only weakly learned or when the word is extremely difficult to see), attention is so centered on the representation of the word itself that the representations of other surrounding words related to that meaning are actively suppressed. Support for this center/surround hypothesis, which is similar to lateral inhibition in perception, was obtained in an additional study by Carr and Dagenbach (1990). Carr and Dagenbach showed that when people attempted to retrieve the meaning of a word that was quickly masked, a large facilitation effect resulted when that priming word itself (the center) was presented as a target. This facilitation was observed even though the same masked prime produced substantial inhibition for related targets (the surround). Thus, these results imply that items in semantic memory may be impaired by an active suppression mechanism when a related retrieval target is sufficiently difficult to retrieve. (It should also be mentioned that such inhibition effects can occur even in the simple task of silently reading a clearly presented priming word when this task is made difficult by virtue of the person having suffered from brain trauma; see, e.g., Blaxton & Bookheimer, 1993; Bushell, in press.) Thus, the results reviewed in this section are congruent with the idea that an item-specific suppression mechanism (akin to that suggested by Postman and Underwood, 1973) can produce retrieval-induced "forgetting" in semantic memory when the retrieval conditions are arduous.

F. Suppression and Retrieval-Induced Forgetting in Episodic Memory

To what degree might retrieval-induced forgetting akin to that observed in semantic memory occur in episodic memory? We have already discussed output interference, in which the retrieval process seemed to cause forgetting (Roediger, 1974). Now we turn to more recent research that specifically examines whether a retrieval-based suppression mechanism might provide a general account of forgetting in episodic memory. Three issues relevant to establishing suppression mechanisms as a cause of long-term episodic forgetting are discussed: (1) the dependence of a target's impairment on its previous tendency to interfere with retrievals of related information; (2) the localization of the impairment to the representation of the previously interfering target item itself; and (3) the durability of episodic retrieval-induced forgetting. Each of these issues has been addressed with the retrieval-practice paradigm (M. C. Anderson, et al., 1994), which we briefly describe next.

1. The Retrieval-Practice Paradigm

The retrieval-practice paradigm was devised as a means of examining the effects of episodic memory retrieval on the subsequent ability to retrieve

competing items (M. C. Anderson et al., 1994). In the retrieval-practice paradigm, people take part in three main phases: a study phase, a retrieval-practice phase, and a final test phase. In the study phase, people study several semantic categories, each composed of several exemplars presented in category–exemplar format (e.g., *fruit–banana*). Next, they perform directed "retrieval practice" on only some of the items that they just studied by completing category-plus-exemplar stem cue tests (e.g., *fruit–ba*____). This retrieval practice is then followed by a final test in which people are cued with each category name and asked to free recall any exemplars of that category that they remember having been presented at any point in the experiment.

Of central interest is the impact of retrieval practice on people's ability to recall the remaining unpracticed exemplars of practiced categories on the delayed-recall test. If retrieving an item renders related items less accessible, performing retrieval practice on some category exemplars should impair recall of the remaining unpracticed exemplars of those categories, relative to the recall of unpracticed exemplars from unpracticed baseline categories. This outcome is indeed what M. C. Anderson et al. found. Whereas performance on practiced items improved rather dramatically relative to the baseline condition (as expected), such facilitation appeared to come at the cost of performance on unpracticed exemplars of practiced categories.

2. Dependence of the Impairment on Concurrent Interference

The impairment observed in the retrieval practice paradigm clearly supports the idea that retrieval causes the forgetting of related episodes. The question arises, however, whether such retrieval-induced forgetting reflects the operation of a suppression mechanisms. For example, the impaired recall performance for unpracticed competitors might instead reflect occlusion (McGeoch, 1936, 1942; Mensink & Raaijmakers, 1988; Raaijmakers & Shiffrin, 1981; Rundus, 1973) from exemplars that were strengthened by retrieval practice. M. C. Anderson et al. (1994) addressed this possibility by examining whether the likelihood of an item suffering retrieval-induced forgetting depended on its tendency to interfere during retrieval practice of its competitors.

According to the suppression account of retrieval-induced forgetting, unpracticed competitors undergo suppression because presenting the shared category cue during retrieval practice leads unpracticed items to interfere during the retrieval of the practice targets. If these assumptions are correct, the more interfering an item can be made, the more likely it will be to suffer from retrieval-induced forgetting. Importantly, and contrary to a competition account, these variations in impairment need not rely on the degree to which practiced items are strengthened. Indeed, if unpracticed competitors are sufficiently noninterfering, no impairment may result at all, even given considerable strengthening of practice competitors.

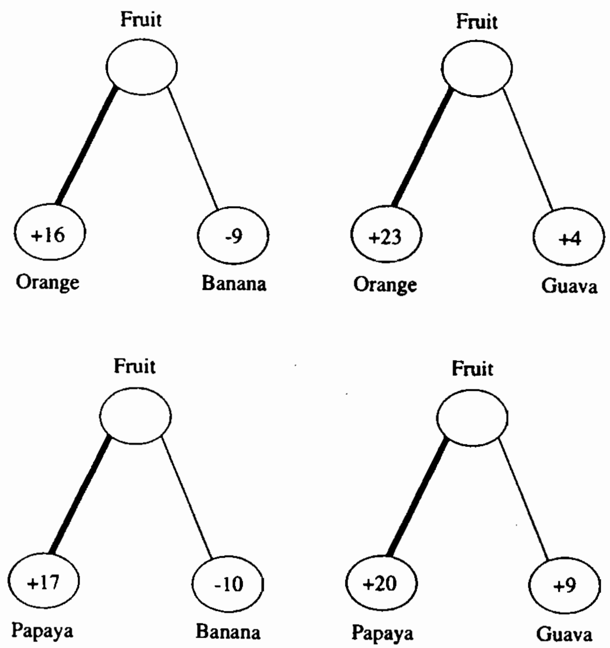

FIGURE 7 Design and results of the retrieval-practice experiment of M. C. Anderson, Bjork, and Bjork (1994) in which they manipulated the taxonomic frequency (high vs. low) of practiced exemplars (left side of each category diagram) and unpracticed competitors (right side of each diagram). Results are depicted in each node as a difference between performance for that item and the corresponding baseline item, with positive numbers designating facilitation and negative numbers designating impairment. The amount of benefit due to retrieval practice may be seen by examining the left side of each node diagram, and the amount of retrieval-induced forgetting on competitors by examining the right side of each diagram.

Figure 7 illustrates some examples of the materials used by M. C. Anderson et al. (1994) to examine the relationship between retrieval-induced forgetting and the degree of interference caused by unpracticed competitors. This figure illustrates four conditions, with the categories and exemplars in each condition depicted as nodes connected by associative links (items that people practiced are depicted on the left of each diagram and unpracticed competitors are depicted on the right). The four conditions differed according to the materials people studied—either high-frequency exemplars (e.g., *orange*) or low-frequency exemplars (e.g., *guava*)—and according to the materials on which people performed retrieval practice. People in both of the top two groups were given retrieval practice on high-frequency members of their categories (e.g., *orange*), but the unpracticed exemplars of their categories were either also high-frequency members

(e.g., *banana*, top left) or were instead low-frequency members (e.g., *guava*, top right). People in both of the bottom two groups were given retrieval practice on low-frequency members of their categories (e.g., *papaya*), but the unpracticed exemplars of their categories were either high-frequency members (e.g., *banana*, lower left) or were low-frequency members (e.g., *guava*, lower right). M. C. Anderson et al. reasoned that if retrieval-induced forgetting depended on the degree to which unpracticed competitors interfered during practice, more retrieval-induced forgetting should occur for high-frequency than for low-frequency exemplars. Furthermore, retrieval-induced forgetting for low-frequency exemplars could be negligible if those items were sufficiently noninterfering. These predictions follow if one assumes high-frequency items to be more interfering than low-frequency items, given the former's stronger pre-experimental associations to the taxonomic category.

Figure 7 also depicts the results of this experiment. Data for each condition appear in the relevant node and are reported here as the difference in percent correct recall between that condition and its corresponding baseline. The positive numbers for practiced items (on the left side of each diagram) indicate the extent to which retrieval practice facilitated recall above baseline performance. The numbers for unpracticed items (on the right side of each diagram) indicate to what extent, if at all, recall of these items was inhibited by retrieval practice on a competitor (with negative numbers indicating inhibition). Unlike the facilitation that occurred for practiced items in all four conditions, the inhibition for unpracticed competitors occurred only when those competitors were high-frequency members of their categories (in upper left and lower left diagrams), with no impairment observed for low-frequency members (upper and lower right diagrams). Thus, consistent with the suppression view, episodic retrieval-induced forgetting occurs only if the target item interfered during the previous retrieval of competitors. Retrieval practice on any set of competitors per se is not sufficient to produce the forgetting, contrary to a pure occlusion account of the effect.

3. Localization of the Impairment to the Item

The greater impact of retrieval practice on high-taxonomic-frequency items favors a suppression-based account over a competition-based account of episodic retrieval-induced forgetting. However, rather than demonstrating suppression, this dependence of impairment on interference might reflect unlearning of the episodic associations linking categories and impaired exemplars, triggered by the intrusion of inappropriate competing responses during retrieval practice. M. C. Anderson and Spellman (1995) addressed this alternative account by examining whether retrieval-induced forgetting is cue independent.

According to the suppression hypothesis, the representations of unpracticed competitors are suppressed during retrieval practice. That is, practicing retrieval of *orange* to the cue *fruit* suppresses the representation of *banana* itself but not the *fruit–banana* association. If the *banana* representation is indeed suppressed, inhibition for *banana* should be measurable when that item is cued by associated retrieval cues other than *fruit*, such as *monkey*. Thus, whereas concurrent interference might be initiated because items share a common retrieval cue, the resulting suppression of interfering items should result in forgetting that generalizes to other cues. M. C. Anderson and Spellman (1995) referred to this predicted property of suppression models as *cue-independent forgetting*, and to the use of a cue such as *monkey* to assess inhibition as the *independent probe method*.

M. C. Anderson and Spellman (1995) examined whether the forgetting observed in the retrieval-practice paradigm exhibited cue independence, using the independent probe method. People followed the retrieval-practice procedure described previously, except that sometimes pairs of practiced and unpracticed categories were related to each other and sometimes they were unrelated. Figure 8 illustrates typical materials from the related and unrelated conditions of one of their experiments. In the related condition (A), people studied the categories *red* and *food*, and performed retrieval practice (indicated by a "+") on some items from one of those categories, for instance, items such as *red–blood*. Though it was never mentioned to the people in the experiment, the nonpracticed items of the *red* study category could also be categorized as *food* and some members of the *food* category could also be categorized as *red*, as shown by the dotted lines linking these items to their related categories. In the unrelated condition (B), categories were not related to one another, replicating the standard retrieval-practice paradigm. The new relations between the practiced and unpracticed categories in the related condition enabled M. C. Anderson and Spellman to test the property of cue independence.

Suppose that impairment of within-category items such as *tomato* results from their suppression when they interfere with *blood* during *red–blood* practice. M. C. Anderson and Spellman (1995) reasoned that if this were so, then *radish* might also be suppressed if, by virtue of its prior semantic association to the *red* category (dotted line), it also interferes with the practice of *blood*. If suppression truly affects item representations, the effects on *radish* should be observable even if *radish* is tested with the independent cue *food*. However, if *tomato*'s impairment results from associative unlearning, the ability to recall *radish* to the cue *food* should remain unaffected, because any impairment caused by practicing *red–blood* should be specific to testing via the unlearned *red–radish* association.

Results of this experiment are depicted in Figure 8 in the nodes for the exemplars in each condition, in terms of the percentage of items correctly

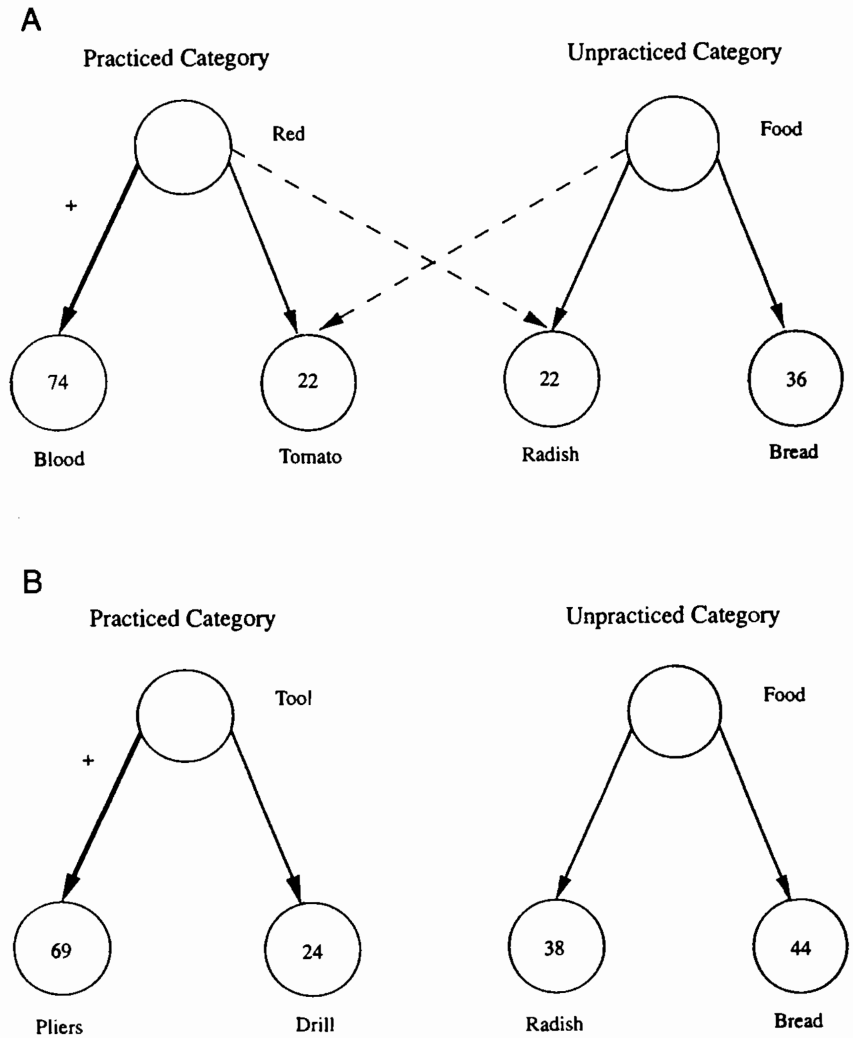

FIGURE 8 Design and results of an experiment by M. C. Anderson and Spellman (1995). Copyright © 1995 by the American Psychological Association. Adapted with permission. (A) The related condition, in which the studied category that was practiced (*red*) and the unpracticed category (*food*) are similar to one another. These categories are similar in that each has members that are categorizable as members of the other category (the figure illustrates this with dotted lines that denote the pre-existing semantic relationships between an item and the other studied category in which it did not appear). (B) The unrelated condition in which the practiced and unpracticed categories are not similar to one another. Note that the items receiving retrieval practice in these two conditions, *blood* and *pliers*, are denoted by a "+" next to the link for each of these items. Results are reported in the nodes for each condition in terms of the percentage of items correctly recalled. The cross-category inhibition of *radish* caused by doing retrieval practice on *blood* may be observed by comparing performance on *radish* in the related condition (A, *top right*) to *radish* in the unrelated condition (B, *bottom right*).

recalled on the final category cued-recall test. The items in the unpracticed category of the unrelated condition (B, *right*) serve as the baseline against which within-category impairment of *tomato* and the between-category impairment of *radish* can be measured in the related condition. As can be seen, retrieval practice on *red–blood* impaired recall of *red–tomato* (A, *left*) on the final recall test, relative to baseline items (i.e., *radish* in the unrelated condition; B, right diagram), replicating the basic retrieval-induced forgetting effect. More important, practicing *red–blood* impaired recall of *radish* to *food* in the related condition, which can be seen by comparing performance for that item to the *radish* baseline in the unrelated condition. Because *radish* was tested under the independent cue *food,* impaired recall of that item cannot be attributed to associative unlearning of the *red–radish* link during the practice of *red–blood*. This cross-category inhibition was observed across three experiments using different materials.

These findings clearly support the view that the retrieval-induced forgetting observed in the retrieval-practice paradigm exhibits the property of cue independence predicted by models attributing impairment to suppression. These findings resemble other phenomenon of cross-cue impairment discussed in previous sections of this chapter, such as output interference and retroactive interference in the A–B, C–D paradigm (both of which also hinge on training or testing procedures requiring retrieval of interfering items). Unlike those previous findings, however, the M. C. Anderson and Spellman (1995) study used a within-subjects baseline against which to measure cross-cue impairment of items like *food–radish*. This difference has the important implication that the present cross-cue impairment may not be explained in terms of competition for a general contextual cue. If participants' ability to recall *food–radish* were impaired because retrieval practice of *red–blood* increased the competition for the experimental context cue, within-subject baseline items should have suffered, as well. As such, there should have been no difference in people's recall performance between the baseline condition and *food–radish* (see M. C. Anderson & Spellman, 1995, for an elaboration of this reasoning). Thus, these findings provide clear evidence for cross-cue impairment that is not predicted by unlearning.

4. Durability of the Impairment

Because the consequences of activation and suppression are generally assumed to be fairly brief (e.g., less than a second), one might wonder whether such processes could underlie long-term forgetting from episodic memory. For example, the retrieval-induced forgetting observed in the retrieval-practice paradigm might reflect impairment that either (1) extends for only a brief period after retrieval practice, or (2) stems entirely from output interference, arising at test, because stronger practiced items tend to

be retrieved before unpracticed competitors. If either of these possibilities characterized episodic retrieval-induced forgetting, such effects could not form the basis of long-term forgetting.

The first possibility—that retrieval-induced forgetting endures for only a brief period after retrieval practice—cannot by itself be true for the simple reason that all of the studies reviewed here have employed retention intervals of 20 min between retrieval practice and the final recall test. We thus turn to the possibility that the impaired recall of unpracticed competitors reflects the fleeting consequences of output interference at the time of test. The contribution of output interference has been assessed by M. C. Anderson et al. (1994), who demonstrated that impairment still occurred on a final recall test when the output order of items within a category was controlled. In this study, unpracticed competitors were always tested prior to stronger practiced competitors through the use of stem cues that uniquely identified each exemplar (e.g., *fruit–a*____) within the experiment. M. C. Anderson and Spellman (1995) observed similar findings, demonstrating that cross-category impairment of items like *food–radish* (see previous section) still occurred on a category-cued recall test when participants were cued to recall the unpracticed category prior to the practiced category. These findings show that retrieval-induced forgetting endures for at least 20 min, rendering it plausible that the mechanisms of retrieval produce the enduring character of episodic recall failure.

G. Retrieval as the Internal Focus of Attention

The findings just reviewed support the notion that retrieval processes in both episodic and semantic memory employ active suppression processes. This suppression causes deficits in the ability to use the affected information on subsequent tasks, which, in the case of episodic memory, may be quite enduring. Thus, such processes might seem quite undesirable. However, several authors (M. C. Anderson et al., 1994; M.C. Anderson & Spellman, 1995; R. A. Bjork, 1989; R. A. Bjork & Bjork, 1992; Dagenbach & Carr, 1994b; Hasher & Zacks, 1988; Keele & Neill, 1978; Neill, 1989; Zacks & Hasher, 1994) have argued for the functional utility of suppression processes as well (see the previous section on directed forgetting for a related discussion). In this section, we consider this functionality in the context of computational models of retrieval in situations that require selective attention. Such considerations have led to the novel claim that retrieval might best be viewed as the internal focus of selective attention, a claim that we hereafter refer to as the *attentional focusing view*. This view suggests a broader approach to understanding why phenomena as diverse as fact retrieval, lexical access, and episodic recollection might share common mechanisms underlying interference.

1. Selective Retrieval and Selective Attention

Most people would agree that attention can be shifted from objects in the external world to objects in the internal world, such as images, facts, or episodes generated on the basis of past experience. For example, to recollect what you had for dinner last evening requires that you cease reading so that attention may be refocused to the mental representation of last evening's dinner. Focusing attention in this manner is likely to bring the contents of that experience into awareness, allowing it to serve as a basis for the requested recollection. When characterized in this manner, the process of accessing particular items of prior knowledge can be viewed as selective attention toward an object that is no longer present in the external world. M. C. Anderson and Spellman (1995) identified these cases of selective retrieval as requiring what they called *conceptually focused selective attention,* thereby highlighting their similarity to cases of perceptually focused selective attention. This similarity is shown in Figure 9, which is taken from M. C. Anderson and Spellman (1995).

Figure 9A depicts the real-world perceptual problem of focusing attention on one piece of fruit in a nearly full fruit bowl. In this figure, the representations of the several fruits receive activation in parallel from per-

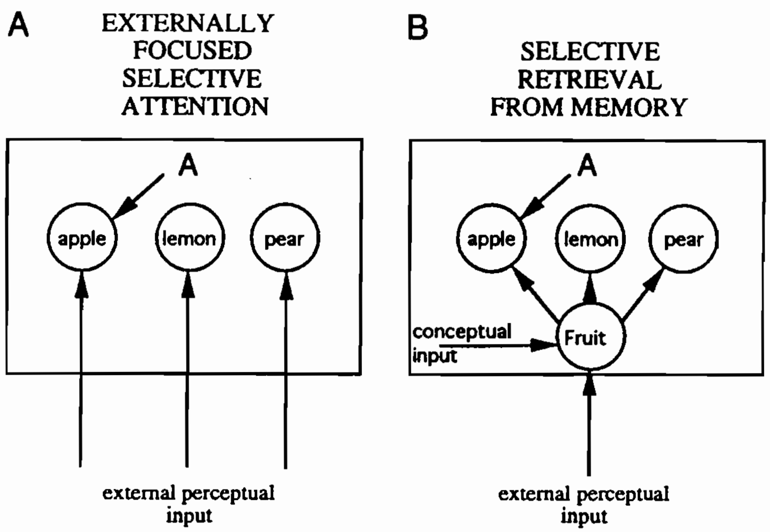

FIGURE 9 A schematic illustration of the relationship between selective attention (A) and selective retrieval from memory (B). In selective attention, attention ("A" in figure) isolates the representation of one fruit (apple) from amongst many fruits activated by external perceptual input. In memory retrieval, attention ("A") isolates the representation of one fruit (apple) from amongst many fruits activated by a retrieval cue (*fruit*), itself activated either by other concepts or by the external world.

ceptual input. Now, consider Figure 9B, which depicts the problem of focusing attention on one fruit among many present in long-term memory. In this figure, the representations of these fruits receive activation in parallel from a shared retrieval cue (which, in turn, may have been activated either by external perceptual input or by activation from other internal conceptual representations). In either case, selectively attending to one of the fruits entails the isolation of its representation from those of its competitors. According to this characterization, the primary differences between selective attention and selective retrieval are: (1) whether the competing representations' activations are conceptually or perceptually initiated; and (2) whether the output of attention is a consciously experienced memory or a consciously experienced percept. In both cases, proper task performance demands the selection of a single representation from among a set of active competitors. If we regard this function of selection as basic to selective attention, then selective retrieval can be regarded as conceptually focused selective attention.

Viewing retrieval as conceptually focused attention has interesting implications for work in both attention and memory. First, the evidence we have reviewed on interference effects in memory would bear on the question of whether attentional selection is achieved via facilitatory or inhibitory processes. Although most theorists agree that attention must somehow facilitate to-be-attended representations, the existence of inhibitory processes for deactivating distracting items is controversial (see Dagenbach & Carr, 1994b, for a collection of papers examining the status of inhibitory processes in cognition). To the extent that memory retrieval can be regarded as conceptually focused attention, data supporting the operation of suppression processes in retrieval and in perceptual selection tasks mutually reinforce each other, suggesting the broader conclusion that suppression achieves attentional selection.

Recent work with the negative-priming paradigm supports such a relation between memory retrieval and perceptually focused selective attention (for excellent reviews of negative priming, see Fox, 1995; Houghton & Tipper, 1994; May, Kane, & Hasher, 1995; Neill, Valdes, & Terry, 1994). In one version of this paradigm (Tipper, 1985), people are presented two line drawings of familiar objects, one in red, the other in green, and are asked to respond to the red one (the target) and to ignore the green one (the distractor). Reaction time to respond to a target on a critical trial, called the probe, is examined as a function of that probe's relation to the immediately preceding trial, called the prime. As might be expected, response times are faster to a target on the probe trial if it had also appeared as a target on the prime trial, relative to a condition in which the two trials are entirely unrelated. More provocative is the finding that response times to the target on the probe trial are longer when that stimulus had been ignored (i.e., was a distractor) on the prime trial. Interestingly, this latter effect, called negative priming, can under some circumstances persist through the presentation of other unrelated items between the prime and probe trials (DeSchepper & Treisman, 1996; Tipper, Weaver, Cameron, Brehaut, & Bestedo, 1991).

The negative-priming paradigm provides a clear test of the situation depicted in Figure 9A, in which attention must be directed to a single item among many activated in a parallel by perceptual input. Thus, the finding that focusing attention to a target percept (i.e., the red item) slows subsequent response times to ignored items (i.e., the green item) supports the hypothesis (Neill, 1977; Tipper, 1985) that attention employs active suppression processes to overcome interference. However, other accounts exist (Lowe, 1979; Park & Kanwisher, 1994). One such account, offered by Neill and Valdes (1992) and Neill, Valdes, Terry, and Gorfein (1992), is analogous to the notion of response competition reviewed in the present chapter. According to this approach, people encode the stimulus that they are supposed to ignore (which we will call the A stimulus) and associate it with the thought "no response," because they understand that they are to respond only to the target item. When that A distractor stimulus appears later as a target, people are assumed spontaneously to retrieve the "no response" feature associated with that A stimulus, which causes competition for the retrieval of the current task-relevant response to A. The extent to which such spontaneous retrievals occur should vary with factors influencing the retrievability of the prime distractor, such as delay and similarity. Although both the suppression and occlusion hypotheses easily accommodate many of the same findings, each accounts for some effects that the other has trouble explaining. Because of this, Fox (1995) and Milliken, Tipper, and Weaver (1994) have argued that both mechanisms are likely to contribute to the negative-priming effect. If these conclusions are correct, the phenomena of negative priming converge with those reviewed in the present chapter, concerning the role of attentional suppression in memory retrieval.

The attentional focusing view of memory retrieval suggests an intriguing perspective on the behavioral causes of forgetting that differs from traditional analyses. Specifically, this view implies that our experiences of retrieval failure may be linked to the very mechanisms that allow us effectively to direct cognition to particular objects in both the external and internal worlds. To understand the implications of this view, consider how semantic memory is used to mediate intelligent behavior, such as reasoning and communication. Because these activities demand prolonged sequential access to very specific referents in memory, they should entail an ongoing barrage of inhibitory suppression that renders access to even well-established conceptual representations volatile. In the case of episodic memory, the elaborate and incidental recollections that occur throughout the day should influence the subsequent accessibility of the episodic representations

activated during those recollections. Thus, whereas classical views of interference have emphasized the role of learning as the primary cause of forgetting, the attentional focusing view emphasizes issues of cognitive control related to the retrieval and use of information, once acquired.

2. Comparisons to Previously Postulated Interference Mechanisms

Although the attentional focusing view of memory retrieval is a relatively new approach to interference effects, many of its aspects resemble components of classical interference theory. We now turn to a discussion of some of the specific relations between the attentional focusing view and three components of classical interference theory: occlusion, unlearning, and response-set suppression. We suggest that the attentional focusing view validates many of the intuitions behind these classical proposals, while at the same time questioning the historical emphasis on associative learning as a cause of forgetting.

a. Occlusion

The occlusion approach assumes that competition, if sufficiently strong, blocks the retrieval of otherwise intact memory items long enough for search efforts to be abandoned. Although the attentional focusing view also assumes that competition impedes recall, the nature of the impediment differs from that in occlusion. In the attentional focusing account, although retrieval competition causes the slowing of recall, it does not directly cause forgetting, as in occlusion; rather, it triggers the need for conceptually focused selective attention. That is, competition among items sharing a cue impedes target retrieval only until inhibitory control processes reduce interference from those competing items, at which point the target will be retrieved.

But how can the attentional focusing view explain the subjective experience of occlusion and its apparent relation to retrieval failure? Such experiences might arise in two ways. First, the retrieval target might be so weakly represented in memory that retrieval failure would occur regardless of whether perseverations of competitors accompanied the experience. To the extent that only partial target information, in the form of a few of the target's featural attributes, is available in memory, items sharing those attributes will tend to be recalled during attempts to retrieve the target. Second, suppression mechanisms might fail to exclude nontarget competitors from consciousness during an initial retrieval, perhaps because of distraction or because those competitors are especially strong. To the extent that retrieval of nontargets entails the (accidental) suppression of the desired item, such intrusions should impair target memory. After being suppressed a first time, the likelihood that suppression will occur again should increase, especially when we consider the increased accessibility of the intruded item. Thus, although the experience of occlusion surely occurs, the attentional focusing view ascribes little role to it as a mechanism of retrieval failure, independent of the suppression that would accompany this experience. Strong evidence for this view comes from M. C. Anderson et al.'s (1994) demonstration (see previous section on retrieval-induced forgetting in episodic memory) that associative strengthening of items via retrieval practice fails to impair competing items if they were unlikely to have been suppressed during that retrieval practice.

b. Unlearning

The conditions and causes of forgetting proposed in the attentional focusing account are quite similar to those proposed in the classical unlearning hypothesis (Melton & Irwin, 1940). According to the unlearning account of retroactive interference in the A–B, A–D paradigm, responses previously learned in List 1 intrude inappropriately during the learning of new responses in List 2. Such intrusions lead to the extinction of the stimulus–response associations linking the A cues to the older intruding B responses, impairing later memory for the A–B pairing. The attentional focusing view emphasizes competitor interference as a prerequisite of impairment, as does unlearning and occlusion.

The unlearning and attentional focusing approaches also both assume that competitor interference triggers a special process by which interfering items are impaired directly, increasing the effectiveness of subsequent retrievals of newly learned information. The attentional focusing account differs from the unlearning proposal, however, in the locus of the impairment and in the mechanisms by which impairment occurs. First, in unlearning, forgetting stems from decrements in cue–target associations, whereas in attentional focusing, impairment is localized to the target item itself. As illustrated in our discussion of the research addressing cue-independent forgetting (M. C. Anderson & Spellman, 1995), impairment associated with memory retrieval generalizes to other cues in a way not predicted by the unlearning hypothesis. Second, the unlearning proposal attributes forgetting to decrements in associative bonds caused by a general learning process, invoked by feedback on whether a retrieved item is appropriate or inappropriate. The attentional focusing view, on the other hand, ascribes forgetting to the reversible suppression of target items caused by an inhibitory control process. This control process reduces interference from competing items, but gives no special importance to the appropriateness of the retrieved item. (That is, suppression of competing items occurs whether a person overtly recalls the target or accidentally intrudes a distractor.) Thus, whereas the present view agrees with the functional conditions proposed in

the unlearning hypothesis, it attributes forgetting to inhibitory control processes rather than to associative learning mechanisms.

c. Response-Set Suppression

Aspects of the attentional focusing view also bear strong resemblance to the response-set suppression hypothesis (Postman et al., 1968). This classical hypothesis attributes a portion of the recall impairment observed in retroactive interference to the suppression of the whole set of intrusive responses from List 1, during acquisition trials on a new set of responses in List 2. According to the response-set suppression hypothesis, the representations of intrusive items are impaired, in reversible fashion, by a control process (i.e., a selector mechanism), functioning to enhance performance in a more current activity. Thus, both the attentional focusing view and response-set suppression emphasize the item representation as the locus of effect and the role of inhibitory control processes in producing impairment.

The attentional focusing view differs from response-set suppression in two respects, however. First, according to Postman et al.'s (1968) classical view, the suppression process is directed at the entire set of List 1 responses and not to any particular item. Although intrusions of particular List 1 items trigger the selector mechanism, the consequent suppression impairs all List 1 items, irrespective of whether they intrude. The attentional focusing account differs in that suppression is directed at individual competitors as a function of whether they compete with target retrieval. Second, the response-set suppression process was construed as a way of shifting response sets, triggered in response to intrusions (see R. A. Bjork, 1989; R. A. Bjork & Bjork, 1992; Zacks & Hasher, 1994, for related functional proposals). In attentional focusing, inhibition occurs as part of the retrieval process, instead of after it, assisting in the isolation of the target item in memory. Nevertheless, even with these differences, the present attentional focusing proposal captures several aspects of response-set suppression, in the context of a theory that considers the on-line dynamics of retrieval.

V. RELATED RESEARCH AREAS

As should be clear from the findings covered in the present chapter, interference effects occur in a broad range of memory paradigms. Because interference is so ubiquitous, space limitations prevent us from covering data from all of the experimental paradigms in which such effects have been studied. Nonetheless, we mention here briefly several areas of research likely to be of interest to many readers, providing references to pertinent sources.

A. Generalizability of Interference Research

The first area concerns whether generalizations about interference as studied with verbal materials apply to other modes of knowledge, and also to other subject populations. Although the present review has focused primarily on interference phenomena in human verbal memory, interference effects in long-term memory have been demonstrated in other knowledge domains, including people's memory for visual stimuli (see, e.g., Chandler, 1989; 1991; Deffenbacher, Carr, & Leu, 1981), for motor skills (Hicks & Cohn, 1975; Hicks & Young, 1973; D. Lewis, McAllister, & Adams, 1951; D. Lewis, Shepard, & Adams, 1949; McAllister & Lewis, 1951; see Adams, 1987, p. 50, for discussion); and for facts accessed during cognitive skills such as mental arithmetic (Campbell, 1987, 1990, 1991; Campbell & Clark, 1989; Campbell, 1995). The basic phenomena and principles of interference studied in human memory also appear to characterize performance in studies of animal learning that employ Pavlovian conditioning paradigms (see Bouton, 1993, for review). Unfortunately, a discussion of how interference effects in these domains relate to those observed in verbal memory is beyond the scope of the present chapter.

Much of the work reported thus far has focused on interference as it occurs in normal populations. However, a great deal of recent work has examined characteristics of interference in special populations in which the management of interference through inhibitory processes may be a crucial issue. For instance, deficits in the ability to inhibit interfering representations have been proposed to underlie cognitive changes associated with normal aging (Gerard, Zacks, Hasher, & Radvansky, 1991; Hartman & Hasher, 1991; Hasher, Stolzfus, Zacks, & Rympa, 1991; Hasher & Zacks, 1988; Kane, Hasher, Stolzfus, Zacks, & Connelly, 1994; McDowd, & Filion, 1992; McDowd, Oseas-Kreger, & Filion, 1994; see also Light, Chapter 13, this volume, for a thorough review of this proposal), and with several clinical syndromes, including schizophrenia (Beech, Powell, McWilliams, & Claridge, 1989; Cohen & Servan-Schreiber, 1992) and frontal lobe dysfunction (Fuster, 1989; Luria, 1966; Mishkin, 1964; Shimamura, 1994). The ability to inhibit extraneous information has even been proposed as an important dimension of general intelligence (Dempster, 1991). The general hypothesis in each of these cases is that the many cognitive deficits observed in these special populations may reflect a more basic deficit in the utilization of attentional inhibition. If these populations suffer from a generalized deficit in the ability to inhibit activated representations, they should exhibit exaggerated susceptibility to interference in both attention and memory tasks. At present, there appears to be promising support for a general, exaggerated susceptibility to interference in these populations, but further work must be done before this characteristic can be confidently attributed to impaired inhibitory processes (see Dagenbach & Carr, 1994b, Dempster & Brainerd, 1994, for collections of reviews of inhibitory processes).

B. The Misinformation Effect in Eyewitness Memory

A second area of research concerns how people's memory reports for the details of a crime event might be affected by encoding misleading information after the original event has been witnessed. Research on this topic began with the classic series of studies by Loftus and colleagues (Loftus, 1975, 1977, 1979a; Loftus, Miller, & Burns, 1978; Loftus & Palmer, 1974) using what has come to be known as the *misinformation procedure*. In a typical version of this procedure, people first view a series of photographic slides depicting a crime event such as a traffic accident or a theft. People then receive additional information about the crime, often in the form of a narrative summarizing the relevant events and details. For participants in the *misled* condition, one of these details is altered, and is thus inaccurate. For example, if the previous slides depicted a car passing a stop sign, the narrative might state that the car had passed a yield sign. People in the *control* condition receive the same narrative, except that the crucial detail is omitted (or in many studies, described in neutral terms, like "sign" instead of "stop sign"). Of crucial interest is whether, on a final recognition test, people select the correct slide on a critical trial in which the originally viewed slide (e.g., of a car passing a stop sign) is paired with a distractor altered to be consistent with the misleading detail (e.g., a slide depicting the same car passing a yield sign).

Numerous studies using this procedure have demonstrated that people receiving misinformation in the intervening narrative select the previously viewed slide significantly less often than do people not receiving misinformation (see Loftus & Hoffman, 1989, for a thorough listing of these studies). For example, in a study by Loftus et al. (1978), misled participants selected the correct slide only 42% of the time, whereas participants in the control condition correctly selected the appropriate slide 75% of the time (numbers estimated from Figure 2 in Loftus et al., 1978). Thus, presenting verbal misinformation after an original event seems substantially to impair people's memory for that event, much like the impairment observed in the classical retroactive interference procedure (compare to Figure 3A). Two aspects of the misinformation effect seem especially striking, especially when considering its relationship to retroactive interference. First, the deficit caused by a single incidental exposure to misinformation appears dramatic: Performance drops from a high level (75%) to worse than would be expected if people were guessing randomly between the two slides (42%). Second, this deficit occurs on a recognition accuracy test, even though retroactive interference is often reduced or even absent on such tests. Based on these considerations and on participants' high confidence when selecting the incorrect slide (Loftus et al., 1978; Tversky & Tuchin, 1989), Loftus has argued that misleading information permanently alters the memory trace for the original event (Loftus, 1975, 1979a, 1979b, 1981; Loftus & Loftus, 1980; Loftus et al., 1978).

Recent work has challenged this memory alteration view, however. Critics argue that misleading information has no effect on people's *memory* for the previously viewed events (McCloskey & Zarragoza, 1985a, 1985b; Zaragoza & Koshmider, 1989; Zaragoza & McCloskey, 1989; Zaragoza, McCloskey, & Jamis, 1987). Rather, decreased overall performance in the misled condition is seen to arise from a select subset of participants in that condition—namely, those participants who, for reasons unrelated to the encoding of misinformation, forgot the initially seen detail and who thus instead selected the altered slide based on their recollection of the misinformation. Although the issues are complex, many authors agree that such response bias or *misinformation acceptance* (Belli, 1989) contributes to the deficit caused by the misinformation procedure (see Belli, 1989; Loftus & Hoffman, 1989). Recent discussion has focused on whether misinformation acceptance is itself an interesting phenomenon (for arguments in favor see Loftus & Hoffman, 1989), whether accessibility is impaired beyond this bias (for supportive evidence see Belli, 1989; Bekerian & Bowers, 1983; Chandler, 1989; Christiannsen & Oschalek, 1983; Tversky & Tuchin, 1989; but see also Zaragoza et al., 1987), and, if so, whether impairment reflects failure to recollect the information in the original event (Chandler, 1991) or confusion over its source (Lindsay, 1990; Lindsay & Johnson, 1988). As discussion of this effect evolves, it is interesting to see that many of the central issues faced by theorists from the classical interference era (e.g., the issue, highlighted in the present chapter, of whether forgetting is produced by unlearning or response competition) have reemerged as pivotal to this debate and its implications for eyewitness memory (see Chandler & Fisher, Chapter 14, this volume, for further discussion of eyewitness memory).

C. Interference on Direct and Indirect Tests of Memory

A third and fairly recent area of research concerns whether interference effects occur on indirect memory tasks. Most of the tasks employed in both classical and modern studies of interference have been direct memory tasks, that is, tasks designed to tap what is now referred to as explicit memory. Direct memory tasks direct participants to make explicit reference to their particular study experience with the prior item to perform the task (e.g., free recall, cued recall, and recognition tests). Over the last decade, considerable attention has been given to how performance on such tasks might differ from performance on tasks that measure memory for a prior episode

without actually asking a person to try to recollect that episode consciously (see Kelley & Lindsay, Chapter 2, this volume; Roediger & McDermott, 1993; Schacter, 1987). Often, research on this topic directly compares performance on indirect and direct tests as follows: People in both conditions study a list of words and then participate in a test of their memory for those words. In a commonly used paradigm, word stems (e.g., *car___*, for *carrot* or *ank__*, for *ankle*) are used as the retrieval cues. In the direct memory test, people are instructed to use each stem as a cue for retrieving a word they had just studied; in the indirect memory test, people are told, without reference to the prior study list, that they should complete the stem with the first word that comes to mind. In both tests, the benefits of prior study are revealed when people are more likely to complete the stem with a studied word than with a nonstudied word.

Considerable data now exist showing that many variables affect performance on direct and indirect tests in different and sometimes opposite ways (see the earlier mentioned reviews). Might these varieties of memory also be differentially susceptible to interference effects of the sort reviewed in the present chapter? The study most relevant to our present coverage was reported by Graf and Schacter (1987), who had people study paired associates in what corresponds to an A–B, A–D paradigm. In Graf and Schacter's experiment, people studied four different new responses to an A stimulus during A–D learning (rather than the customary one). For example, people first studied a pair such as *hen–carrot*, followed later by *hen–pond*, *hen–tree*, *hen–zipper*, and *hen–money*. In the control condition, the stimulus *hen* appeared only in the *hen–carrot* pair. Memory for response members of these paired associates was assessed in two ways, corresponding to the direct and indirect cued-recall methods described previously, except that each stem appeared beside a studied stimulus term (e.g., *hen–car___*). When people were explicitly instructed to retrieve the previously studied item that fit the stem cue, they suffered massive retroactive interference: Stem completion for the studied word was far worse when four other response terms had been studied with that stimulus than when only the cued response itself had been studied. Remarkably, people suffered no retroactive interference in the corresponding indirect memory test! Although a similar dissociation between retroactive interference effects in direct and indirect memory tests has been reported with word lists involving the presentation of individual words rather than paired associates (Jacoby, 1983), other studies using lists of individually presented words have obtained retroactive interference for both indirect and direct tests (see Roediger & McDermott, 1993, for review). One might also question the generality of Graf and Schacter's findings on the basis of the abundance of classical interference findings demonstrating retroactive interference in the MFR test (see previous section on paired associated methodology), which demands only that the subject recall the first verbal response that comes to mind. Nevertheless, the Graf and Schacter (1987) results provide tantalizing, though isolated, evidence that the kinds of interference mechanisms discussed in the present chapter may somehow be short-circuited by indirect memory tests (see, however, related work on directed forgetting in implicit memory: Basden et al., 1993; R. A. Bjork & Bjork, 1991; MacLeod, 1989; Paller, 1990). This intriguing finding clearly merits additional experimental scrutiny.

VI. SUMMARY AND CONCLUSIONS

The function of memory is never so conspicuous and astonishing as when it fails us. One cannot help but wonder how the events of our past—events from the best of times and from the worst of times, events that have shaped our characters and that verify our continued existence—fade to oblivion; or how a concept once well mastered and useful deteriorates into confusion and misunderstanding; or how the name of a friend whom we have known for many years eludes us, even if only momentarily. Such failures abound in daily experience, and sometimes with great consequence, motivating the abundant experimental research that has been devoted to their explanation. In the present chapter, we reviewed research on what many theorists agree is one of the most potent and pervasive factors underlying these experiences: interference.

Interference refers to the impaired ability to remember an item when it is similar to other items stored in memory. Interest in interference was initiated at the turn of the century, when G. E. Mueller and Pilzecker (1900) reported the first empirical study of retroactive interference. This observation ultimately led to a program of research lasting seventy years, the primary focus of which was to explain the substantial forgetting associated with interference. A core advance coming out of this classical era was a simple characterization of the basic situation of interference as competition between items sharing a common retrieval cue. According to this idea, attaching additional memory "responses" to a particular "stimulus" reduces recall performance on target items because those additional items compete with targets upon presentation of their shared stimulus. This core idea permeates modern theoretical work in a variety of domains, although the terms used to characterize competition may vary in each case. Indeed, the last twenty years have illustrated that these elementary dynamics apply to a variety of cognitive tasks in which one must select between multiple concurrently activated mental representations, such as: (1) trying to retrieve a specific episode from our past; (2) trying to retrieve facts that we encoded into semantic memory; (3) performing the seemingly effortless task of retrieving a single word's meaning, spelling, or pronunciation; (4) trying to select, for additional analysis, an object from among other objects in our perceptual environment; and (5) trying to retrieve sequentially the information needed to form a coherent sentence or argument or to solve a problem.

Many theories have been proposed to explain why associating additional items to a retrieval cue might render those items more susceptible to interference. In general, theories addressing this question may be divided into three categories, defined by the locus of the memory representation thought to play the greatest role in causing interference: theories attributing forgetting to changes in (1) the cues that people use to retrieve targets in memory; (2) the associations linking retrieval cues to targets; and (3) the targets themselves. Classical interference factors such as response competition and unlearning (comprising the Two-Factor Theory, Melton & Irwin, 1940) emphasized the role of associative learning processes, fitting our experience of forgetting into the general mold of learning theories prevalent during that era. Late in the development of the classical interference period, additional factors such as variable-stimulus encoding (E. Martin, 1968, 1971) and response-set suppression (Postman et al., 1968) were proposed, in which forgetting was attributed to changes in the retrieval cue and the memory target representations, respectively. Unfortunately, although these later theories of interference have considerable plausibility, they have not received the attention given to associative interference theories, in part, because the focus of memory research shifted from interference shortly after their development (see Wheeler, 1995, for a similar perspective). Since the end of the classical interference era, theories of interference have focused primarily on mechanisms of response competition (specifically, occlusion), reflecting a general disenchantment with unlearning and a belief that interference is best conceived as retrieval inhibition, that is, the failure to retrieve otherwise available memory items.

In the final section of this chapter, we presented a recent perspective on the causes of memory interference that casts our experiences of forgetting in a very different light. According to this perspective, our tendency to forget is intimately linked with the very mechanisms that allow focused memory retrieval to occur. That is, forgetting of target items derives not from learning interfering information per se, but from the selective retrieval of that interfering information after it has been acquired. Selectively retrieving related items harms our later recall of critical target memories by means of an active suppression process that inhibits those critical targets; although this suppression process helps to overcome competition exerted by these critical items, it has the side effect of impairing the retrieval of those targets when they later become relevant. Because inhibitory control processes are thought to be triggered in situations much like those we face in situations requiring selective attention, it is argued that retrieval should be regarded as a case of conceptually focused selective attention. We argued herein that this view validates many of the intuitions underlying classical theories of forgetting, while questioning the widespread assumption that forgetting derives in any direct way from associative learning. Rather, our many losses—of memories of past experiences, of our friends' names, or of our comprehension of concepts with which we once were adept—are seen as costs of the very mechanisms that enable us to direct cognition to internal thoughts and to the external environment.

Acknowledgments

The authors would like to thank Mike Ciranni, Chad Dodson, and Arthur Shimamura for commenting on early drafts of this chapter, and Elizabeth Bjork, Bob Bjork, C. C. Chandler, and Henry Roediger for help with the final version. Writing of this paper was supported by postdoctoral Grant 94-21 to M. C. Anderson by the McDonnell-Pew foundation, and by Grant MH48757 to Arthur Shimamura from the National Institute of Mental Health.

References

Adams, J. A. (1987). Historical review and appraisal of research on the learning, retention and transfer of human motor skills. *Psychological Bulletin, 101,* 41–74.

Anderson, J. R. (1974). Retrieval of propositional information from long-term memory *Cognitive Psychology, 6,* 451–474.

Anderson, J. R. (1976). *Language, memory and thought.* Hillsdale, NJ: Erlbaum.

Anderson, J. R. (1983). *The architecture of cognition.* Cambridge: Harvard University Press.

Anderson, M. C., & Bjork, R. A. (1994). Mechanisms of inhibition in long-term memory: A new taxonomy. In D. Dagenbach & T. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 265–325). San Diego: Academic Press.

Anderson, M. C., Bjork, R. A., & Bjork, E. L. (1994). Remembering can cause forgetting: Retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 1063–1087.

Anderson, M. C., & Spellman, B. A. (1995). On the status of inhibitory mechanisms in cognition: Memory retrieval as a model case. *Psychological Review, 102,* 68–100.

Anderson, R. C., & Watts, G. H. (1971). Response competition in the forgetting of paired associates. *Journal of Verbal Learning and Verbal Behavior, 10,* 29–34.

Balota, D. A., Ferraro, F. R., & Connor, L. T. (1991). On the early influence of meaning on word recognition: A review of the literature. In P. Schwanenflugel (Ed.), *The psychology of word meanings* (pp. 187–222). Hillsdale, NJ: Erlbaum.

Barnes, J. M., & Underwood, B. J. (1959). "Fate" of first-list associations in transfer theory. *Journal of Experimental Psychology, 58,* 95–105.

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.

Beech, A., Powell, T., McWilliams, J., & Claridge, G. (1989). Evidence of reduced "cognitive inhibition" in schizophrenia. *British Journal of Clinical Psychology, 28,* 110–116.

Bekerian, D. A., & Bowers, J. M. (1983). Eyewitness testimony: Were we misled? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9,* 139–145.

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

Bjork, E. L., Bjork, R. A., & Glenberg, A. (1973, November). *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., & 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* (pp. 217–235). Washington, DC: Winston.

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

Bjork, R. A., & Bjork, E. L. (1991). *Dissociations in the impact of to-be-forgotten information on memory.* Paper presented at the 99th Annual Convention of the American Psychological Association, San Francisco.

Bjork, R. A., & Bjork, E. L. (1992). A new theory of disuse and an old theory of stimulus fluctuation. In A. F. Healy, S. M. Kosslyn, & R. M. Shiffrin (Eds.), *Essays in honor of William K. Estes: Vol. 1. From learning theory to connectionist theory* (pp. 35–67). Hillsdale, NJ: Erlbaum.

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

Blaxton, T. A., & Bookheimer, S. Y. (1993). Retrieval inhibition in anomia. *Brain and Language, 22,* 221–237.

Blaxton, T. A., & Neely, J. H. (1983). Inhibition from semantically related primes: Evidence of a category-specific retrieval inhibition. *Memory & Cognition, 11,* 500–510.

Bouton, M. E. (1993). Context, time and memory retrieval in the interference paradigms of Pavlovian learning. *Psychological Bulletin, 114,* 80–99.

Bower, G. H., Thompson-Schill, S., & Tulving, E. (1994). Reducing retroactive interference: An interference analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 51–66.

Brown, A. S. (1979). Priming effects in semantic memory retrieval processes. *Journal of Experimental Psychology: Human Learning and Memory, 5,* 65–77.

Brown, A. S. (1981). Inhibition in cued retrieval. *Journal of Experimental Psychology: Human Learning and Memory, 7,* 204–215.

Brown, A. S. (1991). The tip of the tongue experience: A review and evaluation. *Psychological Bulletin, 109,* 204–223.

Brown, J. (1968). Reciprocal facilitation and impairment in free recall. *Psychonomic Science, 10,* 41–42.

Brown, R., & McNeill, D. (1966). The tip-of-the-tongue phenomenon. *Journal of Verbal Learning and Verbal Behavior, 5,* 325–337.

Burgess, C., & Simpson, G. B. (1988). Cerebral hemispheric mechanisms in the retrieval of ambiguous word meanings. *Brain and Language, 33,* 86–103.

Bushell, C. (in press). Dissociated identity and semantic priming in Broca's aphasia: How controlled processing produces inhibitory semantic priming. *Brain and Language.*

Campbell, J. I. D. (1987). Network interference and mental multiplication. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13,* 109–123.

Campbell, J. I. D. (1990). Retrieval inhibition and interference in cognitive arithmetic. *Canadian Journal of Psychology, 44,* 445–464.

Campbell, J. I. D. (1991). Conditions of error priming in number-fact retrieval. *Memory & Cognition, 19,* 197–209.

Campbell, J. I. D. (1995). Mechanisms of simple addition and multiplication: A modified network-interference theory and stimulation. *Mathematical Cognition, 1,* 121–164.

Campbell, J. I. D., & Clark, J. M. (1989). Time course of error priming in number fact retrieval: Evidence for excitatory and inhibitory mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 920–929.

Carr, T. H., & Dagenbach, D. (1990). Semantic priming and repetition priming from masked words: Evidence for a center-surround attentional mechanism in perceptual recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 341–350.

Ceraso, J., & Henderson, A. (1965). Unavailability and associative loss in RI and PI. *Journal of Experimental Psychology, 70,* 300–303.

Ceraso, J., & Henderson, A. (1966). Unavailability and associative loss in RI and PI: Second try. *Journal of Experimental Psychology, 72,* 314–316.

Chandler, C. C. (1989). Specific retroactive interference in modified recognition tasks: Evidence for an unknown cause of interference. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 341–350.

Chandler, C. C. (1991). How memory for an event is influenced by related events: Interference in modified recognition tests. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 115–125.

Chandler, C. C., & Gargano, G. J. (1995). Item-specific interference caused by cue-dependent forgetting. *Memory & Cognition, 23,* 701–708.

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

Coe, W. C., Basden, B. H., Basden, D., & Fikes, T. (1989). Directed forgetting and posthypnotic amnesia: Information processing and social contexts. *Journal of Personality and Social Psychology, 56,* 189–198.

Cohen, J. D., & Servan-Schreiber, D. (1992). Context, cortex and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review, 99,* 45–77.

Conway, A. R. A., & Engle, R. W. (1994). Working memory and retrieval: A resource dependent inhibition model. *Journal of Experimental Psychology: General, 123,* 354–373.

Crowder, R. G. (1976). *Principles of learning and memory.* Hillsdale, NJ: Erlbaum.

Dagenbach, D., & Carr, T. H. (1994a). Inhibitory processes in perceptual recognition: Evidence for a center-surround attentional mechanism. In D. Dagenbach & T. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 327–358). San Diego: Academic Press.

Dagenbach, D., & Carr, T. H. (Eds.). (1994b). *Inhibitory processes in attention, memory, and language.* San Diego: Academic Press.

Dagenbach, D., Carr, T. H., & Barnhardt, T. M. (1990). Inhibitory semantic priming of lexical decisions due to failure to retrieve weakly activated codes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 328–340.

Dagenbach, D., Carr, T. H., & Wilhelmson, A. (1989). Task-induced strategies and near-threshold priming: Conscious influences on unconscious perception. *Journal of Memory and Language,* 412–443.

Davis, J. C., & Okada, R. (1971). Recognition and recall of positively forgotten items. *Journal of Experimental Psychology, 89,* 181–186.

Deffenbacher, K. A., Carr, T. H., & Leu, J. R. (1981). Memory for words, pictures, and faces: Retroactive interference, forgetting, and reminiscence. *Journal of Experimental Psychology: Human Learning and Memory, 7,* 299–305.

Dempster, F. N. (1991). Inhibitory processes: A neglected dimension of intelligence. *Intelligence, 15,* 157–173.

Dempster, F. N., & Brainerd, C. J. (Eds.). (1994). *New perspectives on interference and inhibition.* New York: Academic Press.

DeSchepper, B., & Treisman, A. M. (1996). Visual memory for novel shapes: Implicit coding without attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 27–47.

Epstein, W. (1972). Mechanisms of directed forgetting. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 6, pp. 147–191). New York: Academic Press.

Epstein, W., Massaro, D. W., & Wilder, L. (1972). Selective search in directed forgetting. *Journal of Experimental Psychology, 94,* 18–24.

Epstein, W., & Wilder, L. (1972). Searching for to-be-forgotten material in a directed forgetting task. *Journal of Experimental Psychology, 95,* 349–357.

Erdelyi, M. H. (1985). *Psychoanalysis: Freud's cognitive psychology.* New York: W. H. Freeman.

Erdelyi, M. H. (1993). Repression: The mechanism and the defense. In D. M. Wegner & J. W. Pennebaker (Eds.), *Handbook of mental control* (pp. 126–148). Englewood Cliffs, NJ: Prentice-Hall.

Erdelyi, M. H., & Goldberg, B. (1978). Let's not sweep repression under the rug: Towards a cognitive psychology of repression. In J. F. Kihlstrom & F. J. Evans (Eds.), *Functional disorders of memory* (pp. 355–402). Hillsdale, NJ: Erlbaum.

Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review, 62,* 369–377.

Estes, W. K. (1972). An associative basis for coding and organization in memory. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory* (pp. 161–191). New York: Wiley.

Fox, E. (1995). Negative priming from ignored distractors in visual selection: A review. *Psychonomic Bulletin & Review, 2,* 145–173.

Fuster, J. M. (1989). *The prefrontal cortex: Anatomy, physiology, and neuropsychology of the frontal lobe.* New York: Raven Press.

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. L. (1983). Disrupted retrieval in directed forgetting: A link with posthypnotic amnesia. *Journal of Experimental Psychology: General, 112,* 58–72.

Geiselman, R. E., MacKinnon, D. P., Fishman, D. L., Jaenicke, C., Larner, B., Shoenberg, S., & Swartz, S. (1983). Mechanisms of hypnotic and non-hypnotic forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9,* 626–635.

Gelfand, H., & Bjork, R. A. (1985, November). *On the locus of retrieval inhibition in directed forgetting.* Paper presented at the meeting of the Psychonomic Society, Boston.

Gerard, L., Zacks, R. T., Hasher, L., & Radvansky, G. A. (1991). Age deficits in retrieval: The fan effect. *Journal of Gerontology, 46,* 131–136.

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review, 91,* 1–67.

Graf, P., & Schacter, D. L. (1987). Selective effects of interference on implicit and explicit memory for new associations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13,* 45–53.

Greenberg, S., & Wickens, D. D. (1972). Is matching performance an adequate test of "extinction" effects on individual associations? *Psychonomic Science, 27,* 227–229.

Hartman, M., & Hasher, L. (1991). Aging and suppression: Memory for previously relevant information. *Psychology and Aging, 6,* 587–594.

Hasher, L., Stolzfus, E. R., Zacks, R. T., & Rympa, B. (1991). Age and inhibition. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 17, 163–169.

Hasher, L., & Zacks, R. T. (1988). Working memory, comprehension and aging: A review and a new view. In G. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 22, 193–225). San Diego: Academic Press.

Hicks, R. E., & Cohn, D. M. (1975). Lack of proactive inhibition in a psychomotor task at two different retention intervals. *Journal of Motor Behavior,* 7, 101–104.

Hicks, R. E., & Young, R. K. (1973). Effect of amount of interpolated learning and length of

retention interval upon retroactive inhibition in a serial search task. *Journal of Experimental Psychology, 100,* 297–301.
Holmes, D. S. (1990). The evidence for repression: An examination of sixty years of research. In J. L. Singer (Ed.), *Repression and dissociation: Defense mechanisms and personality styles* (pp. 85–102). Chicago: University of Chicago Press.
Horton, K. D., & Petruk, R. (1980). Set differentiation and depth of processing in the directed forgetting paradigm. *Journal of Experimental Psychology: Human Learning and Memory, 6,* 599–610.
Houghton, G., & Tipper, S. P. (1994). A model of inhibitory mechanisms in selective attention. In D. Dagenbach & T. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 53–112). San Diego: Academic Press.
Jacoby, L. L. (1983). Perceptual enhancement: Persistent effects of an experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9,* 21–83.
Johnson, H. M. (1994). Processes of successful intentional forgetting. *Psychological Bulletin, 116,* 274–292.
Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness.* Cambridge: Harvard University Press.
Jongeward, R. H., Woodward, A. E., & Bjork, R. A. (1975). The relative roles of input and output mechanisms in directed forgetting. *Memory & Cognition, 3,* 51–57.
Joordens, S., & Besner, D. (1994). When banking on meaning is not (yet) money in the bank: Explorations in connectionistic modeling. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 1051–1062.
Kane, M. J., Hasher, L., Stolzfus, E. R., Zacks, R. T., & Connelly, S. L. (1994). Inhibitory attentional mechanisms and aging. *Psychology and Aging, 9,* 103–112.
Karchmer, M. A., & Winograd, E. (1971). Effects of studying a subset of familiar items on recall of the remaining items: The John Brown effect. *Psychonomic Science, 25,* 224–225.
Keele, S. W., & Neill, W. T. (1978). Mechanisms of attention. In E. C. Carterette & P. Friedman (Eds.), *Handbook of perception IX* (pp. 3–47). New York: Academic Press.
Kihlstrom, J. F. (1977). Models of posthypnotic amnesia. In W. E. Edmonston (Ed.), *Conceptual and investigative approaches to hypnosis and hypnotic phenomena* (Vol. 296, pp. 284–301). New York: New York Academy of Sciences.
Kihlstrom, J. F. (1980). Posthypnotic amnesia for recently learned material: Interactions with "episodic" and "semantic" memory. *Cognitive Psychology, 12,* 227–251.
Kihlstrom, J. F. (1983). Instructed forgetting: Hypnotic and non-hypnotic. *Journal of Experimental Psychology: General, 112,* 73–79.
Kihlstrom, J. F., & Barnhardt, T. M. (1993). The self-regulation of memory: For better and for worse, with and without hypnosis. In D. M. Wegner & J. W. Pennebaker (Eds.), *Handbook of mental control* (pp. 88–125). Englewood Cliffs, NJ: Prentice-Hall.
Kihlstrom, J. F., & Evans, F. J. (1979). Memory retrieval processes in posthypnotic amnesia. In J. F. Kihlstrom & F. J. Evans (Eds.), *Functional disorders of memory* (pp. 179–218). Hillsdale, NJ: Erlbaum.
Lewis, C. H., & Anderson, J. R. (1976). Interference with real world knowledge. *Cognitive Psychology, 8,* 311–335.
Lewis, D., McAllister, D. E., & Adams, J. A. (1951). Facilitation and interference in performance on the Modified Mashburn Apparatus: I. The effects of varying the amount of original learning. *Journal of Experimental Psychology, 41,* 247–260.
Lewis, D., Shephard, A. H., & Adams, J. A. (1949). Evidence of associative interference in psychomotor performance. *Science, 110,* 271–273.
Lewis, M. Q. (1971). Categorized lists and cued recall. *Journal of Experimental Psychology, 87,* 129–131.
Lindsay, D. S. (1990). Misleading suggestions can impair eyewitnesses' ability to remember

event details. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 1077–1083.
Lindsay, D. S., & Johnson, M. K. (1988). The eyewitness suggestibility effect and memory for source. *Memory & Cognition, 17,* 349–358.
Loftus, E. F. (1975). Leading questions and the eyewitness report. *Cognitive Psychology,* 7, 560–572.
Loftus, E. F. (1977). Shifting human color memory. *Memory & Cognition, 5,* 696–699.
Loftus, E. F. (1979a). *Eyewitness testimony.* Cambridge: Harvard University Press.
Loftus, E. F. (1979b). The malleability of memory. *American Scientist, 67,* 312–320.
Loftus, E. F. (1981). Mentalmorphosis: Alterations in memory produced by the mental bonding of new information to old. In J. Long & A. Baddeley (Eds.), *Attention and performance IX* (pp. 417–434). Hillsdale, NJ: Erlbaum.
Loftus, E. F., & Hoffman, H. G. (1989). Misinformation in memory: The creation of new memories. *Journal of Experimental Psychology: General, 118,* 100–104.
Loftus, G. R., & Loftus, E. F. (1974). The influence of one memory retrieval on a subsequent memory retrieval. *Memory & Cognition, 2,* 467–471.
Loftus, E. F., & Loftus, G. R. (1980). On the permanence of stored information in the human brain. *American Psychologist, 35,* 409–420.
Loftus, E. F., Miller, D. G., & Burns, H. J. (1978). Semantic integration of verbal information into a visual memory. *Journal of Experimental Psychology: Human Learning and Memory, 4,* 19–31.
Loftus, E. F., & Palmer, J. E. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior, 13,* 585–589.
Lowe, D. G. (1979). Strategies, context, and the mechanism of response inhibition. *Memory & Cognition,* 7, 392–389.
Luce, R. D. (1959). *Individual choice behavior.* New York: Wiley.
Luria, A. R. (1966). Disturbances of higher cortical functions with lesions of the frontal regions. In *Higher cortical functions in man* (pp. 218–295). London: Tavistock.
MacLeod, C. M. (1975). Long-term recognition and recall following directed forgetting. *Journal of Experimental Psychology: Human Learning and Memory, 1,* 271–279.
MacLeod, C. M. (1989). Directed forgetting affects both direct and indirect tests of memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 13–21.
Marcel, A. J. (1980). Conscious and preconscious recognition of polysemous words: Locating the selective effects of prior verbal context. In R. S. Nickerson (Ed.), *Attention and performance VIII* (pp. 435–457). Hillsdale, NJ: Erlbaum.
Martin, D. W., & Kelly, R. T. (1974). Secondary task performance during directed forgetting. *Journal of Experimental Psychology, 103,* 1074–1079.
Martin, E. (1968). Stimulus meaningfulness and paired-associate transfer: An encoding variability hypothesis. *Psychological Review, 75,* 421–441.
Martin, E. (1971). Verbal learning theory and independent retrieval phenomena. *Psychological Review, 78,* 314–332.
May, C. P., Kane, M. J., & Hasher, L. (1995). Determinants of negative priming. *Psychological Bulletin, 118,* 35–54.
McAllister, D. E., & Lewis, D. (1951). Facilitation and interference in performance on the Modified Mashburn Apparatus: II. The effects of varying the amount of interpolated learning. *Journal of Experimental Psychology, 41,* 356–363.
McCloskey, M., & Bigler, K. (1980). Focused memory search in fact retrieval. *Memory & Cognition, 8,* 253–264.
McCloskey, M., & Zaragoza, M. (1985a). Misleading postevent information and memory for events: Arguments and evidence against memory impairment hypotheses. *Journal of Experimental Psychology: General, 114,* 1–16.

McCloskey, M., & Zaragoza, M. (1985b). Postevent information and memory: Reply to Loftus, Schooler, and Wagenaar. *Journal of Experimental Psychology: General, 114*, 381–387.
McDowd, J. M., & Filion, D. L. (1992). Aging, selective attention, and inhibitory processes: A psychophysiological approach. *Psychology and Aging, 7*, 65–71.
McDowd, J. M., Oseas-Kreger, D. M., & Filion, D. L. (1994). Inhibitory processes in selective attention and aging. In F. N. Dempster & C. J. Brainerd (Eds.), *New perspectives on interference and inhibition in cognition* (pp. 363–400). New York: Academic Press.
McGeoch, J. A. (1936). Studies in retroactive inhibition: VII. Retroactive inhibition as a function of the length and frequency of presentation of the interpolated lists. *Journal of Experimental Psychology, 19*, 674–693.
McGeoch, J. A. (1942). *The psychology of human learning*. New York: Longmans, Green.
McGovern, J. B. (1964). Extinction of associations in four transfer paradigms. *Psychological Monographs, 78*(16, Whole No. 593).
Melton, A. W., & Irwin, J. M. (1940). The influence of degree of interpolated learning on retroactive inhibition and the overt transfer of specific responses. *American Journal of Psychology, 3*, 173–203.
Melton, A. W., & von Lackum, W. J. (1941). Retroactive and proactive inhibition in retention: Evidence for a two-factor theory of retroactive inhibition. *American Journal of Psychology, 54*, 157–173.
Mensink, G. J. M., & Raajmakers, J. G. W. (1988). A model of interference and forgetting. *Psychological Review, 95*, 434–455.
Merryman, C. T. (1971). Retroactive inhibition in the A–B, A–D paradigm as measured by a multiple-choice test. *Journal of Experimental Psychology, 91*, 212–214.
Merryman, C. T., & Merryman, S. S. (1971). Stimulus encoding in the A–B′, AX–B and the A–B′r, AX–B paradigms, *Journal of Verbal Learning and Verbal Behavior, 10*, 681–685.
Meyer, A. S., & Bock, K. (1992). The tip-of-the-tongue phenomenon: Blocking or partial activation? *Memory & Cognition, 20*, 715–726.
Milliken, B., Tipper, S. P., & Weaver, B. (1994). Negative priming in a spatial localization task: Feature mismatching and inhibition of distractor location. *Journal of Experimental Psychology: Human Perception and Performance, 20*, 624–646.
Mishkin, M. (1964). Perseveration of central sets after frontal lesions in monkeys. In J. M. Warren & K. Akert (Eds.), *The frontal granular cortex and behavior* (pp. 219–237). New York: McGraw-Hill.
Mueller, C. W., & Watkins, M. J. (1977). Inhibition from part-set cueing: A cue-overload interpretation. *Journal of Verbal Learning and Verbal Behavior, 16*, 699–710.
Mueller, G. E., & Pilzecker, A. (1900). Experimentalle Beitrage zur Lehre vom Gedachtnis. *Zeitschrift fur Psychologie, 1*, 1–300.
Myers, J. L., O'Brien, E. J., Balota, D. A., & Toyofuku, M. L. (1984). Memory search without interference: The role of integration. *Cognitive Psychology, 16*, 217–242.
Neely, J. H. (1976). Semantic priming and retrieval from lexical memory: Evidence for facilitatory and inhibitory processes. *Memory & Cognition, 4*, 648–654.
Neely, J. H., Schmidt, S. R., & Roediger, H. L. (1983). Inhibition from related primes in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9*, 196–211.
Neill, W. T. (1977). Inhibitory and facilitatory processes in selective attention. *Journal of Experimental Psychology: Human Perception and Performance, 3*, 444–450.
Neill, W. T. (1989). Lexical ambiguity and context: An activation-suppression model. In D. S. Gorfein (Ed.), *Resolving semantic ambiguity* (pp. 63–83). New York: Springer-Verlag.
Neill, W. T., & Valdes, L. A. (1992). Persistence of negative priming: Steady state or decay? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18*, 565–576.

Neill, W. T., Valdes, L. A., & Terry, K. M. (1994). Selective attention and the inhibitory control of cognition. In F. N. Dempster & C. J. Brainerd (Eds.), *New perspectives on interference and inhibition in cognition* (pp. 207–261). New York: Academic Press.
Neill, W. T., Valdes, L. A., Terry, K. M., & Gorfein, D. S. (1992). Persistence of negative priming: II. Evidence for episodic trace retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18*, 993–1000.
Nickerson, R. S. (1984). Retrieval inhibition from part-set cueing: A persisting enigma in memory research. *Memory & Cognition, 12*, 531–552.
Paller, K. A. (1990). Recall and stem-completion priming have different electrophysiological correlates and are modified differentially by directed forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*, 1021–1032.
Park, J., & Kanwisher, N. (1994). Negative priming for spatial location: Identity mismatching, not distractor inhibition. *Journal of Experimental Psychology: Human Perception and Performance, 20*, 613–623.
Postman, L. (1971). Transfer, interference and forgetting. In J. W. Kling & L. A. Riggs (Eds.), *Woodworth and Schlosberg's: Experimental psychology* (3rd ed., pp. 1019–1132). New York: Holt, Rinehart & Winston.
Postman, L., & Stark, K. (1969). The role of response availability in transfer and interference. *Journal of Experimental Psychology, 79*, 168–177.
Postman, L., Stark, K., & Fraser, J. (1968). Temporal changes in interference. *Journal of Verbal Learning and Verbal Behavior, 7*, 672–694.
Postman, L., & Underwood, B. J. (1973). Critical issues in interference theory. *Memory & Cognition, 1*, 19–40.
Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review, 88*, 93–134.
Radvansky, G. A., & Zacks, R. T. (1991). Mental models and the fan effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17*, 940–953.
Ratcliff, R., Clark, S. E., & Shiffrin, R. M. (1990). The list-strength effect: I. Data and discussion. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*, 163–178.
Ratcliff, R., & Murdock, B. B., Jr. (1976). Retrieval processes in recognition memory. *Psychological Review, 83*, 190–214.
Reason, J. T., & Lucas, D. (1984). Using cognitive diaries to investigate naturally occurring memory blocks. In J. E. Harris & P. E. Morris (Eds.), *Everyday memory actions and absentmindedness* (pp. 53–70). London: Academic Press.
Reder, L. M. (1982). Plausibility judgments vs. fact retrieval: Alternative strategies for sentence verification. *Psychological Review, 89*, 250–280.
Reder, L. M., & Anderson, J. R. (1980). A partial resolution of the paradox of interference: The role of integrating knowledge. *Cognitive Psychology, 12*, 447–472.
Reder, L. M., & Ross, B. H. (1983). Integrated knowledge in different tasks: The role of retrieval strategy on fan effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9*, 55–72.
Reder, L. M., & Wible, C. (1984). Strategy use in question-answering: Memory strength and task constraints on fan effects. *Memory & Cognition, 12*, 411–419.
Reitman, W., Malin, J. T., Bjork, R. A., & Higman, B. (1973). Strategy control and directed forgetting. *Journal of Verbal Learning and Verbal Behavior, 12*, 140–149.
Richardson, J., & Stanton, S. K. (1972). Some effects of learning to a set of components on stimulus selection. *American Journal of Psychology, 85*, 519–533.
Roediger, H. L. (1973). Inhibition in recall from cueing with recall targets. *Journal of Verbal Learning and Verbal Behavior, 12*, 644–657.
Roediger, H. L. (1974). Inhibiting effects of recall. *Memory & Cognition, 2*, 261–269.
Roediger, H. L. (1978). Recall as a self-limiting process. *Memory & Cognition, 6*, 54–63.

Roediger, H. L., & McDermott, K. B. (1993). Implicit memory in normal human subjects. In F. Boller & J. Grafman (Eds.), *Handbook of neuropsychology* (Vol. 8, pp. 63–131). New York: Elsevier.

Roediger, H. L., & Neely, J. H. (1982). Retrieval blocks in episodic and semantic memory. *Canadian Journal of Psychology, 36,* 213–242.

Roediger, H. L., Neely, J. H., & Blaxon, T. (1983). Inhibition from related primes in semantic memory retrieval: A reappraisal of Brown's (1979) paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9,* 478–489.

Roediger, H. L., & Schmidt, S. R. (1980). Output interference in the recall of categorized and paired associate lists. *Journal of Experimental Psychology: Human Learning and Memory, 6,* 91–105.

Roediger, H. L., Stellon, C., & Tulving, E. (1977). Inhibition from part-list cues and rates of recall. *Journal of Experimental Psychology: Human Learning and Memory, 3,* 174–188.

Roediger, H. L., & Tulving, E. (1979). Exclusion of learned material from recall as a post-retrieval operation. *Journal of Verbal Learning and Verbal Behavior, 18,* 601–615.

Rudy, J. W. (1974). Stimulus selection in animal conditioning and paired-associate learning: Variations in the associative process. *Journal of Verbal Learning and Verbal Behavior, 13,* 282–296.

Rundus, D. (1973). Negative effects of using list items as recall cues. *Journal of Verbal Learning and Verbal Behavior, 12,* 43–50.

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

Schneider, N. G., & Houston, J. P. (1968). Stimulus selection and retroactive inhibition. *Journal of Experimental Psychology, 77,* 166–167.

Shimamura, A. P. (1994). Memory and frontal lobe function. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 803–814). Cambridge: MIT Press.

Simpson, G. B., & Kang, H. (1994). Inhibitory processes in the recognition of homograph meanings. In D. Dagenbach & T. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 359–382). San Diego: Academic Press.

Slamecka, N. J. (1968). An examination of trace storage in free recall. *Journal of Experimental Psychology, 76,* 504–513.

Slamecka, N. J. (1972). The question of associative growth in the learning of categorized materials. *Journal of Verbal Learning and Verbal Behavior, 11,* 324–332.

Slamecka, N. J. (1975). Intralist cueing of recognition. *Journal of Verbal Learning and Verbal Behavior, 14,* 630–637.

Smith, A. D. (1971). Output interference and organized recall from long-term memory. *Journal of Verbal learning and Verbal Behavior, 10,* 400–408.

Smith, E. E., Adams, N., & Schorr, D. (1978). Fact retrieval and the paradox of interference. *Cognitive Psychology, 10,* 438–464.

Springer, S. P., & Deutsch, G. (1993). *Left brain, right brain* (4th ed.). New York: W. H. Freeman.

Timmins, W. K. (1974). Varying processing time in directed forgetting. *Journal of Verbal Learning and Verbal Behavior, 13,* 539–544.

Tipper, S. P. (1985). The negative priming effect: Inhibitory priming and ignored objects. *Quarterly Journal of Experimental Psychology, 37A,* 571–590.

Tipper, S. P., Weaver, B., Cameron, S., Brehaut, J. C., & Bastedo, J. (1991). Inhibitory mechanisms of attention in identification and localization tasks: Time course and disruption. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 681–692.

Todres, A. K., & Watkins, M. J. (1981). A part-set cueing effect in recognition memory. *Journal of Experimental Psychology: Human Learning and Memory, 7,* 91–99.

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of Memory* (pp. 381–403). New York: Academic Press.

Tulving, E., & Arbuckle, T. Y. (1963). Sources of intratrial interferences in paired-associate learning. *Journal of Verbal Learning and Verbal Behavior,* 1, 321–334.

Tulving, E., & Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review, 80,* 352–373.

Tversky, B., & Tuchin, M. (1989). A reconciliation of the evidence on eyewitness testimony: Comments on McCloskey and Zaragoza (1985). *Journal of Experimental Psychology: General, 118,* 86–91.

Tzeng, O. J., Lee, A. T., & Wetzel, C. D. (1979). Temporal coding in verbal information processing. *Journal of Experimental Psychology: Human Learning and Memory, 5,* 52–64.

Underwood, B. J. (1948). 'Spontaneous recovery' of verbal associations. *Journal of Experimental Psychology, 38,* 429–439.

Watkins, M. J. (1975). Inhibition in recall with extralist "cues." *Journal of Verbal Learning and Verbal Behavior, 14,* 294–303.

Watkins, M. J. (1978). Engrams as cuegrams and forgetting as cue-overload: A cueing approach to the structure of memory. In C. R. Puff (Ed.), *The structure of memory* (pp. 347–372). New York: Academic Press.

Wegner, D. M. (1994). Ironic processes of mental control. *Psychological Review, 101,* 34–52.

Wegner, D. M., & Pennebaker, J. W. (1993). *Handbook of mental control.* Englewood Cliffs, NJ: Prentice-Hall.

Weiner, B. (1968). Motivated forgetting and the study of repression. *Journal of Personality, 36,* 213–234.

Weiner, B., & Reed, H. (1969). Effects of instructional sets to remember and to forget on short-term retention: Studies of rehearsal control and retrieval inhibition (repression). *Journal of Experimental Psychology, 79,* 226–232.

Wetzel, C. D., & Hunt, R. E. (1977). Cue delay and the role of rehearsal in directed forgetting. *Journal of Experimental Psychology: Human Learning and Memory, 3,* 233–245.

Wheeler, M. A. (1995). Improvement in recall over time without repeated testing: Spontaneous recovery revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 173–184.

Whitlow, J. W. (1984). Effects of precuing on focused search in fact retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* 733–744.

Wickelgren, W. A. (1976). Memory storage dynamics. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes: Vol. 4: Attention and memory* (pp. 321–362). Hillsdale, NJ: Erlbaum.

Williams, R. F., & Underwood, B. J. (1970). Encoding variability: Test of the Martin hypothesis. *Journal of Experimental Psychology, 86,* 317–324.

Woodward, A. D., & Bjork, R. A. (1971). Forgetting and remembering in free recall: Intentional and unintentional. *Journal of Experimental Psychology, 89,* 109–116.

Zacks, R. T., & Hasher, L. (1994). Directed ignoring: Inhibitory regulation of working memory. In D. Dagenbach & T. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 241–264). San Diego: Academic Press.

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.

Zaragoza, M. S., & McCloskey, M. (1989). Misleading postevent information and the memory impairment hypothesis: Comment on Belli and reply to Tversky and Tuchin. *Journal of Experimental Psychology: General, 118,* 92–99.

Zaragoza, M. S., McCloskey, M., & Jamis, M. (1987). Misleading postevent information and recall of the original event: Further evidence against the memory impairment hypothesis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13,* 36–44.