2

Retrieval as a Self-limiting Process: Part II

ROBERT A. BJORK, ELIZABETH L. BJORK, and
BETHANY J. CAUGHEY

Across much of the history of systematic research on human memory, three stages critical to the memory process have been postulated: trace formation, trace storage, and trace utilization or retrieval (see, e.g., Melton, 1963). It was not, however, until Tulving and Pearlstone's (1966) seminal study that the fundamental importance of retrieval processes in human memory—and the need to distinguish the conditions affecting storage from those affecting retrieval—became clear. By providing cues to participants during the recall of items from categorized lists, Tulving and Pearlstone were able to demonstrate that what is stored or available in memory cannot be indexed in a precise or reliable way by what is recallable or accessible at a given point in time.

Prior to the Tulving and Pearlstone (1966) study, a typical paradigm used by researchers to investigate both storage and retrieval processes was to ask participants to recall all the items they could following the presentation of different types of lists, and it was simply assumed that what was recalled reflected what information had been successfully stored and, furthermore, how that information was organized in memory. Items on a list that are semantically related tend to be recalled together, for example, even if they are presented at disparate positions in the list, which was taken as evidence that interitem associations formed at the time of storage then guided the retrieval process at the time of recall. In other words, the retrieval process underlying recall was thought to be self-propagating: The recall of one item triggers the recall of a next, associated item, which triggers the recall of yet another item, and so forth. Such a self-propagating property was—and still is, to a large extent—intrinsic to associative and organizational theories of the recall process.

In J. R. Anderson's (1972) model of recall, for example, it was assumed that to-be-learned items, such as a list of randomly presented words, are associated with each other at the time of study and that a participant's later recall of the list is guided and sustained by such interitem associations. In other words, the items recalled first serve as cues for items to be recalled later. As Roediger (1978) pointed out, this self-propagating nature of recall—that is, the act of recall itself providing the cues that guide additional recall—was originally proposed by Aristotle and "the intervening 2,000 years have done little to damage the idea" (p. 54).

During about the same period, however—that is, the late 1960s and early 1970s—the results of other studies demonstrated that the act of recall could also function to *limit* further recall. One such example is the research of Slamecka (1968, 1969) demonstrating inhibition owing to part-list cueing. If, following the presentation of a list of words to be remembered, some items from the list are given as cues for the remaining items in the list, the recall of the remaining items is not enhanced, but, rather, is impaired. The fact that inhibition owing to part-list cueing has been referred to as an "enigma for memory researchers" (Nickerson, 1984) reflects not just that it has proved difficult to understand from a detailed process standpoint, but also that it is highly counterintuitive, given that the self-propagating nature of recall seems so compelling.

A second example of the self-limiting property of recall comes from research on the recall of categorized lists. If, following the presentation of a list consisting of multiple exemplars of a number of categories, participants are then cued with category names to aid their recall of members of a given category, recall of such category members decreases with the output position of the provided category cue. That is, the later in the recall test that a given category name is presented as a cue, the fewer members of that category are recalled, a demonstration of "output interference" (e.g., Smith, 1971; Smith, D'Agostino, & Reid, 1970).

ROEDIGER'S ROLE: RETRIEVAL AS A SELF-LIMITING PROCESS

Henry L. Roediger III, inspired by such effects, became one of the first researchers to champion the notion that the act of recall can function to attenuate further recall. He embarked on a program of research during the 1970s and early 80s that revealed the robustness of such negative effects, which he frequently referred to as *recall interference* (e.g., Roediger & Schmidt, 1980). In what follows, we briefly review Roediger's body of work, focusing on the studies that examined the recall of categorized lists—studies that led Roediger to conceptualize recall as a self-limiting process and to argue for hierarchical theories of memory structure (e.g., Estes, 1972; Mandler, 1967; Rundus, 1973) versus interitem associative theories of memory, such as Anderson's (1972) model. We then discuss a new type of recall interference—the more recently discovered phenomenon of *retrieval-induced forgetting*—and we conclude with a speculation that a suppression mechanism that has been put forward to explain retrieval-induced forgetting (e.g., Anderson, Bjork, & Bjork, 1994; Anderson & Spellman, 1995) might also, when elaborated, account for earlier retrieval-interference effects as well.

Positive and Negative Consequences of Retrieval Cues

Beginning with his work as a graduate student in Robert Crowder's laboratory at Yale University, Roediger carried out numerous studies demonstrating that retrieval is a "self-limiting process" (1978). From a theoretical standpoint, he struggled with why retrieval had both positive and negative consequences, depending on the particular paradigm and procedure employed in a given experiment. Research carried out by Endel Tulving and his collaborators and by Norman Slamecka and his collaborators, as summarized below, was particularly influential from that standpoint.

As demonstrated by Tulving and his colleagues (e.g., Tulving & Osler, 1968; Tulving & Pearlstone, 1966), the amount of information that can be recalled from memory depends critically on the nature of the retrieval cues provided at the time of recall. After presentation of a blocked categorized list, for example, participants cued with category names recall significantly more items from the list than do participants simply asked to free recall as many items as they can remember from the list. As argued by Tulving and Pearlstone, such results indicate both that more information is stored and available in memory than can be accessed on a typical free recall test and that the presentation of words in a categorized list apparently determines both the way that they are then organized in memory and later retrieved from memory.

If, as indicated by these results, the items presented in a list to be memorized are stored in such a dependent manner (e.g., in terms of direct associative bonds among the items), then it should be possible to reveal such dependence by presenting some of the items from the list as retrieval cues to aid participants' recall of the remainder of the items in the list. Reasoning in just this way, Slamecka (1968, 1969) conducted a series of studies presenting subsets of items from the presented list as retrieval cues for the rest of the list. To his great surprise, providing some list items not only did not facilitate recall of the remaining words, they appeared instead—over a wide variation of the proportion of list items presented as cues—to have a slight but consistently *negative* effect on the recall of the remaining items. Even when using categorized lists, as in the Tulving and Pearlstone study (1966), Slamecka found presentation of items as retrieval cues to impair, rather than facilitate, recall of the remaining items. The following quotation summarizes Slamecka's reaction to his own findings.

> In the face of our compelling preconceived expectations, we were annoyingly disappointed to find that the experiment failed to show any advantage for the cued group, and worse, that it revealed a small but significant inhibitory effect in that condition. . . . [Six experiments later], I had no choice but to conclude that the classical theoretical portrayal of memory traces as being joined by direct associative links was wrong.
>
> (Slamecka, 1984, p. 96)

Interpreting the Effects of Retrieval Cues

Given these apparently inconsistent effects of providing retrieval cues—that is, the facilitative effects found by Tulving and Pearlstone (1966) versus the deleterious effects found by Slamecka (1968, 1969), Roediger (1973, 1974) argued that a major challenge to the understanding of memory was to determine the conditions under which the provision of retrieval cues benefit versus impair recall and, further, to provide a compelling theoretical explanation to account for these differing effects of retrieval cues. As a step toward this latter goal, Roediger (1973, 1974) proposed a two-factor explanation for when retrieval cues would be beneficial and when they would be detrimental. Basically, he argued that the effects are positive when cues enable access to more higher-order units (e.g., the names of categories) than can be recalled by a participant unaided by such cues, and the effects are negative when more retrieval cues than are needed to access higher-order units are provided (e.g., additional instances from the category).

The reasoning behind Roediger's proposed second factor was that the impairment resulting from the presentation of list cues on recall of the remaining list items could be an instance of output interference—the deleterious effects of earlier recall upon later recall, as first systematically studied by Tulving and Arbuckle (1963). Although Tulving and Arbuckle, on the basis of their work with short paired-associate lists, argued that the phenomenon of output interference was limited to recall of only very recently presented material—that is, recall of items from primary memory—subsequent research with longer lists and delayed recall tests demonstrated that output interference occurred in recall from secondary/long-term memory as well (e.g., Dong, 1972; Smith 1971; Smith et al., 1970). If, as Roediger (1973) proposed, the presentation of cues to participants could be assumed to simulate their own uncued retrieval of those items, then the presentation of these cues should have the same effect as the participants having retrieved these items early in their own recall output; namely, they would impair the recall of later items.

To test the adequacy of this second factor for predicting when providing retrieval cues would impair additional recall, Roediger (1973) performed an experiment in which—following the presentation of fairly long and blocked categorized lists—participants were given either category names as retrieval cues or, in some cases, category names plus a varying number of instances from the categories. Roediger found impaired recall of additional category instances when both category names and instances were provided as retrieval cues, versus when only category names were provided as retrieval cues, with the detrimental effect increasing as the number of instances given as cues increased—a pattern that was consistent with Roediger's two-factor hypothesis as to when retrieval cues would and would not be beneficial.

Additionally, owing to the construction of his categorized list, Roediger (1973) was able to rule out a plausible alternative explanation of these findings involving guessing. By showing that the detrimental effect persisted even when the remaining number of items available for recall was held constant, he was able to demonstrate that the increasingly detrimental effect on recall as the number of provided cues increased did not stem from the fact that as the number of cues increase the number of yet-to-be-recalled items typically decrease, making participants—in principle, at least—more likely to output a correct response by guessing when fewer cues are presented.

Roediger's (1973) study thus went a long way toward resolving what had appeared to be contradictory findings regarding the effects of retrieval cues on recall. When retrieval cues increased access to higher-order units, they increased participant's recall, consistent with the findings of Tulving and Pearlstone (1966) and Tulving and Osler (1968). When category instances were also provided as cues, however, recall of additional items was impaired, with the impairment increasing as the number of instances provided as cues increased, consistent with the part-list cueing effects observed by Slamecka (1968, 1969).

Roediger (1973), like Slamecka (1968, 1969) before him, concluded that his findings were inconsistent with interitem associative theories of memory. He argued instead for hierarchical theories of memory and, in particular, that a model proposed by Rundus (1973) to account for output interference effects in recall could be generalized to account for the observed negative effects of providing retrieval cues as well. Rundus's model, which provided a kind of framework for Roediger's research on recall as a self-limiting process (e.g., Roediger, 1973, 1978; Roediger & Schmidt, 1980; Roediger, Stellon, & Tulving, 1977), is described in the next section.

Rundus's (1973) Model of Recall

The Rundus (1973) model of recall, which builds upon Shiffrin's (1970) model of memory search, was proposed to account for the retrieval process underlying recall and, in particular, to account for the occurrence of output-interference effects in recall. In his model, Rundus assumed that during the presentation of a list of items to be learned, the learner (or participant) attempts to organize the list in some way, and whatever ideas or units the participant uses for such organization then become the higher-order retrieval cues (Tulving, 1966) for the subset of list items organized under it. In the case of a categorized list (e.g., one containing instances of fruits, flowers, birds, etc.), the names of the categories are assumed to serve as these higher-order organizing units. The nature of the hierarchical organizational structure that would result from the presentation of a list, based on the Rundus model, is illustrated in Figure 2.1.

As indicated in Figure 2.1, the words (W) presented in the list are grouped into subsets according to their association with the higher-order retrieval cues (RQ), which in the case of a categorized list would be the names of the categories, and the strength of association between a given higher-order retrieval cue, say RQ_i, and an item under it, say W_j, is denoted as A_{ij}. These higher-order RQs are to be thought of as control elements (Estes, 1972) in the sense that they control the recallability of the words under them. Further, the RQs themselves are assumed to be associated with a contextual cue (denoted as "List" in Figure 2.1) specifying the particular list under consideration. In other words, the association of all the RQs to this contextual cue would indicate that the categories all occurred in that particular

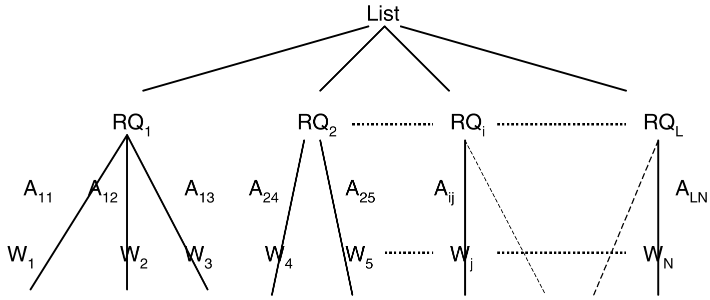

FIGURE 2.1 Organization in memory of items (W) and retrieval cues (RQ) after presentation of a list as proposed by the Rundus model of recall (based on Rundus, 1973).

list. It is also assumed that the strength of association between the various RQs and the list control element, as well as that between a given RQ and the items in its subset, will not be equal, perhaps owing to different amounts of covert rehearsal during list presentation or owing to pre-experimentally established associations in the case of categorized lists.

Three rules assumed in the model determine how the process of recall unfolds. First, the likelihood that a given RQ is retrieved when the participant is asked to recall the list is determined by a ratio rule. Specifically, the likelihood of a given RQ being retrieved is equal to the ratio of the strength of that RQ's association to the "List" cue to the sum of the strengths of association of all RQs to the "List" cue. After retrieval of one of the RQs, the participant attempts to retrieve the items subsumed under that retrieval cue, and the probability of recalling a particular item is determined by the same ratio rule (i.e., the strength of association of that item to that particular RQ divided by the sum of strengths of association of all items to that RQ). If the retrieved item has not been recalled yet, the participant outputs that item and then returns to the same RQ to continue the recall process. If the item had been output previously, that fact would be noted before returning to the RQ, but, critically, it is assumed that the retrieval process involves sampling with replacement, such that previously retrieved words or categories remain available for future recall attempts. Thus, according to this second rule—sampling with replacement—RQs and items (e.g., category names and instances, respectively) may be re-retrieved. Also of critical importance to the model's ability to account for output interference is the assumption that the act of recalling a unit (either the name of a category or an instance) increases the strength of its association to its higher-order control element.

Finally, the model incorporates a stopping rule. According to this rule, attempts to retrieve items associated to a given RQ (e.g., instances within a given category) will continue until a series of k consecutive retrievals produces no new items, at which point the participant abandons further use of that RQ, returns to the context "List" cue to retrieve another RQ, and begins retrieving items organized under that RQ. The participant is assumed to continue in this manner until a series of m consecutively retrieved RQs produces no new items for recall. Thus, in the model, the parameter k determines when the participant decides that continuing to search for new items under the current RQ is unlikely to be productive and, consequently, samples another RQ; and the parameter m determines when the participant decides that further sampling of RQs is unlikely to produce any new items for recall and, consequently, stops the recall process for that list altogether.

Given these assumptions, the Rundus (1973) model can account for both output interference and part-list cueing effects. Output interference is predicted because recall of an item strengthens the association between that item and its RQ, making it more likely to be retrieved again owing to the ratio rule and, at the same time, making it less likely that some other item associated with the same RQ will be retrieved. Additionally, recall of an item increases the likelihood that a participant will stop searching for additional items under that RQ and will go on to sample a new RQ owing to the stopping rule. That is, as more items are retrieved from the same RQ, the probability increases that the criterion number k of no new items retrieved is reached, causing the participant to abandon further retrieval attempts from that RQ. The negative consequences of part-list cueing can be explained in terms of the model by assuming that—when participants are given a subset of the list just studied as cues for recall of the remaining items—it is actually the presented items that tend to be retrieved first with these retrievals acting to block access to the other items from the list.

Tests of Rundus's Model

Over the next 10 years or so, against a theoretical background that pitted hierarchical models against horizontal-associative models, Roediger and colleagues continued to investigate the self-limiting nature of recall via empirical studies on the dynamics of output interference and inhibition owing to part-list cueing. In one particularly important study, Roediger et al. (1977)—by increasing the time available to participants to recall and then measuring the rate of recall—were able to demonstrate that the detrimental effect of cues is not overcome by allowing participants to spend a longer time attempting to recall. Additionally, by comparing the performance of participants who were allowed to output the list items presented as cues as well as any remaining items in the list (i.e., target items) with the performance of participants who were only allowed to output the target items, Roediger et al. were able to rule out the notion that the negative effects of part-list cues might arise from the necessity for participants to check each retrieved item against the set of provided cues. Although relieving participants of this extra checking task by allowing them to output cues as well as target items did increase their recall of target items when compared to the condition where they were not allowed to recall cues, it did not eliminate the negative effects of part-list cueing relative to free recall.

In another series of experiments, Roediger (1978) also obtained retrieval-interference effects consistent with hierarchical associative theories of recall. When he provided participants with some but not all category names as cues, recall for cued categories was facilitated while recall of noncued categories was impaired relative to free recall (Experiment 1) and the impaired recall of noncued categories increased as the number of cued categories from the list increased (Experiment 2). In a third experiment, in which recall of noncued categories was delayed by various types of interpolated activities, he was able to rule out that the impaired recall of noncued categories relative to free recall was attributable to nonspecific effects of the prior recall of cued categories—analogous, perhaps to the way that any interpolated task, such as counting backwards by threes, might impair recall.

On the basis of these and other findings, Roediger (1978) argued that the assumptions of Rundus's (1973) model allow for *both* the self-propagating and self-limiting characteristics of recall. Recall of items within the cued categories is aided, relative to free recall, by providing the participant with the RQs or control elements under which those items were hierarchically or vertically associated. Because, however, the RQs or control elements for the cued categories are strengthened by such recall, access to the noncued RQs or control elements is decreased, owing to both the ratio rule for determining probability of retrieval and the assumption that sampling is with replacement. Furthermore, consistent with the assumptions of the Rundus model, the consequence of presenting cues for some categories but not others is to impair the recall of the noncued categories, not the items within the categories. Finally, the detrimental effect of providing *both* category names and instances is explained by the Rundus model in the same manner as part-list cueing effects are explained—namely, by assuming that the presentation of the additional instances has the same effects as does covert recall of those items by the participant and, thus, represents an instance of output interference. In terms of the Rundus model, the instances presented as cues are strengthened, making them more likely to be retrieved again owing to the sampling-with-replacement rule and thereby to block retrieval of any additional items from that category.

In summary, although some of the results obtained by Roediger and others (see, in particular, Roediger & Schmidt, 1980; Roediger et al., 1977; and Watkins, 1975) could not be explained—at least not efficiently, by Rundus's model—the bulk of Roediger's findings from his many cueing experiments proved consistent with the model. To this day, in fact, Rundus's (1973) model—at the conceptual level, if not at the detailed mathematical-model level—remains the dominant explanation of output interference and inhibition owing to part-list cueing.

We turn now to a discussion of the retrieval-practice paradigm and to retrieval-induced forgetting, which constitutes a more recent type of evidence that retrieval is a self-limiting process.

RETRIEVAL-INDUCED FORGETTING

Retrieval-induced forgetting (RIF), first reported by Anderson et al. (1994), refers to the negative impact on the recall of some items associated to a given cue or configuration of cues when other items associated with that cue are repeatedly retrieved. The practiced items become more recallable than they would have been without such practice, which is hardly surprising, and the unpracticed items become less recallable than corresponding unpracticed items from unpracticed categories—which *is* surprising.

The retrieval-practice paradigm, as implemented by Anderson et al. (1994), has its roots in research on test effects—or, said differently, on "retrieval as a memory modifier" (Bjork, 1975), another research domain in which Henry Roediger has been, and continues to be, heavily involved. The specific procedure used by Anderson et al. was developed as a means of assessing the effects of increasing the retrieval strength of some items on the retrieval strength of related items—and, more specifically, to test some predictions of a "new theory of disuse" proposed by Bjork and Bjork (1992) in a chapter honoring one of the greatest theorists in the history of research on human memory, William K. Estes. The theory, designed to account for a number of "important peculiarities" of human memory, consists of a set of assumptions as to how study and retrieval events impact two presumed dimensions of an item's representation in memory—its "storage strength" and its "retrieval strength"—as a function of that item's current levels of those strengths.

The key assumptions of the theory for present purposes, as summarized by Bjork (2001), are the following.

- Memory representations are double indexed in memory—by their current "retrieval strength" (how accessible or active they are) and their "storage strength" (how well learned or interassociated they are with other memory representations). Storage strength is assumed to accumulate as a consequence of study or practice and, once accumulated, is permanent. Retrieval strength, however, which completely determines the probability of being able to access a given stored representation, is volatile. It is assumed to increase as a consequence of study or practice, but to decrease as a consequence of study or practice of competing responses or behaviors. The theory is a "new" theory of disuse because, in contrast to Thorndike's (1914) original "law of disuse," it is access to learned representations (retrieval strength) that is lost over a period of disuse, not the representation *per se* (storage strength).
- In distinguishing between storage strength and retrieval strength, the theory resurrects a distinction that was common among learning theorists of an earlier era. The distinction is essentially the same, for example, as Hull's (1943) distinction between *habit strength* and *momentary excitatory potential*, or Estes's (1955) distinction between *habit strength* and *response strength*. The distinction also corresponds, in a general way, to the time-honored distinction between learning and performance, a distinction necessitated by a wide range of findings from research on both humans and animals: What we observe is *performance*; what we are often trying to infer is *learning*. Storage strength and retrieval strength also correspond, roughly, to Tulving's distinction between the *availability* and *accessibility* of memory representations (see, e.g., Tulving & Pearlstone, 1966).
- What *is* new about the theory are the assumptions governing how the

current storage and retrieval strengths of a representation influence (a) the increments in the storage strength of that representation that result from study or practice; and (b) the increments and decrements, respectively, in the retrieval strength of that representation that result from study or practice of that representation or competing representations. The assumptions of special pertinence to the issues of the present chapter are the following:

(1) Storage strength serves to enhance the gain and retard the loss of retrieval strength. That is, access to representations in memory, as indexed by retrieval strength, is lost more slowly with disuse—and regained more rapidly given study or practice—the higher that representation's current storage strength.
(2) The higher the current retrieval strength of a representation, the smaller the increments in both storage strength and retrieval strength that result from study or practice of that representation. Thus, somewhat surprisingly, the more accessible a representation, the smaller the increment in storage strength (learning) that results from additional study or practice of that representation. Put differently, conditions that result in forgetting (loss of retrieval strength) also create opportunities for additional learning (i.e., increments in storage strength).

From the foregoing assumptions, retrieval as a self-limiting process emerges in an intrinsic way—because retrieving some item associated to a given cue not only increases that item's retrieval strength, but also decreases the retrieval strength of other items associated with that cue. The theory also makes the counterintuitive prediction that items high in retrieval strength—other things being equal—will be the most susceptible to retrieval-induced forgetting.

The RIF Paradigm and Basic Findings

As implemented by Anderson et al. (1994), the RIF paradigm includes four distinct phases: a study phase, a directed retrieval-practice phase, a distractor phase, and a final test phase. In the study phase, participants are given a list of materials categorized in some way and the retrieval-practice phase involves directed retrieval of some members of some of the categories. In Anderson et al.'s initial experiments, participants were presented with six members of each of eight categories (e.g., *Fruit Orange*; *Weapon Rifle*), with the pairs presented one pair at a time and intermixed by category. The retrieval-practice phase consisted of repeated recall of half the members of half the categories, triggered by cues such as *Fruit–Or*________. To increase the effectiveness of such retrieval practice, each practiced pair was given three such tests separated by expanding intervals (cf. Landauer & Bjork, 1978). Finally, after a retention interval (20 minutes) filled with a distracting activity, participants were asked to try to recall all of the members of all of the categories. Depending on the experiment, the final test consisted of either category-cued free recall or—as a means of controlling the order of output of the members of a given category—cued recall consisting of prompts such as *Fruit B*________, *Fruit O*________ for the six members of a given category.

Of central interest in the RIF paradigm are the levels of recall of three different types of items: (a) practiced exemplars from practiced categories (e.g., *Orange*), called *Rp+* items; (b) unpracticed exemplars from practiced categories (e.g., *Banana*), called *Rp–* items); and (c) unpracticed exemplars from unpracticed categories (e.g., *Rifle*), called *Nrp* or baseline items. The typical pattern of findings is that the recall of Rp+ items exceeds the recall of the Nrp or baseline items, consistent with a long history of research on test effects (e.g., Allen, Mahler, & Estes, 1969; Bjork, 1975; Landauer & Bjork, 1978; Roediger & Karpicke, 2006; Whitten & Bjork, 1977), and, of considerably more interest, however, that the recall of the Rp– items is *impaired* relative to the recall of the Nrp or baseline items. Such retrieval-induced forgetting of Rp– items is surprising because one might expect such items also to profit from the retrieval-practice phase, owing to dynamics such as covert rehearsal or spreading activation of the type intrinsic to horizontal interitem associative models of memory.

The pattern of results observed by Anderson et al. (1994) has now been replicated many times, not only with category-exemplar pairs, but also with a wide range of other materials, including newly-learned visuospatial materials (Ciranni & Shimmamura, 1999); visual scenes and event narratives in eye-witness memory (e.g., Saunders & MacLeod, 2002; Shaw, Bjork, & Handal, 1995); examination materials (Macrae & MacLeod, 1999), autobiographical memories (Barnier, Hung, & Conway, 2004); and stereotypical and/or valenced attributes of hypothetical individuals (Dunn & Spellman, 2003; Storm, Bjork, & Bjork, 2005).

Interpreting Retrieval-Induced Forgetting

In principle, RIF could also reflect blocking and sampling-with-replacement dynamics of the type favored by Roediger and others—in the context of the Rundus (1973) model—as an explanation of output interference—and, more generally, retrieval as a self-limiting process. That is, it could be the case that when participants try to recall all the members of a practiced category, the Rp+ items come readily and repeatedly to mind, blocking access to the Rp– items and impairing their recall.

For several reasons, however, Anderson et al. (1994) argued that they could reject blocking-type explanations of RIF. One such reason is that RIF was still obtained when the order of recall of items from a given category was controlled. That is, even when Rp– items were the first to-be-recalled items from a given practiced category (in response to prompts such as *Fruit–B*________), their recall was impaired relative to corresponding Nrp items. A second reason is that Anderson et al. found that it was the strong exemplars of a given category, such as *Orange* or *Banana*, that were most susceptible to RIF; whereas, the recall of weak exemplars, such as *Guava* and *Papaya*, was not impaired by the practice of other exemplars. Such a finding is consistent with the predictions of Bjork and Bjork's (1992) new theory of disuse, but is inconsistent with blocking as instantiated by ratio-rule mechanisms, which predicts the opposite pattern (for a more detailed version of that argument, see Anderson et al., 1994, Appendix A, p. 1085).

As an alternative explanation of RIF, Anderson et al. (1994) argued for an inhibition/suppression mechanism. The basic idea is that responding to a retrieval-practice prompt such as *Fruit–Or*________ requires not only selecting *Orange*, but also suppressing other exemplars of the *Fruit* category. The fact that it is the strong exemplars of a given category that are most susceptible to RIF is consistent with such a mechanism—because such exemplars are the most likely to come to mind in response to *Fruit* and, hence, are the most frequent targets of suppression. Recently, in research on the possible role of RIF in impression formation, Storm, Bjork, and Bjork (2005) have also found that it is the more recallable attributes of a given hypothetical person, positive or negative, that are most susceptible to RIF. Once again, the argument is that such attributes are the ones that have to be selected against most often, leading to their suppression. It needs to be stressed, however, that this suppression is assumed to occur automatically and without conscious attention; that is, as some information is retrieved, other competing information is selected against and thereby suppressed, with its suppression occurring rather like a by-product of the selection process.

In a second study, Anderson, Bjork, and Bjork (2000) found that the direction of retrieval practice also mattered with respect to whether RIF was obtained, and mattered in a fashion consistent with such a suppression mechanism. When practice was prompted in the standard direction, recalling *Orange* in response to *Fruit–Or*________, RIF was obtained; when practice was prompted in the other direction—that is, recalling *Fruit* in response to *Fr*________*–Orange*—there was no RIF of unpracticed members of the *Fruit* category. Such a result is consistent with a suppression mechanism because it is only retrieval practice in the standard direction that requires the selection of *Orange* and the suppression of other studied fruits.

The fact that the direction of retrieval practice mattered not only supports the suppression idea, but also—in combination with the fact that the strengthening effects of retrieval practice did not depend on direction—argues against a blocking mechanism. That is, Anderson et al. (2000) found that the subsequent recall of Rp+ items profited to the same degree, relative to Nrp items, whether retrieval practice was in the *Fr*________*–Orange* direction or in the *Fruit–Or*________ direction. Given that finding, the blocking of access to Rp– items by Rp+ items should also be commensurate in the two conditions, according to any unelaborated implementation of the blocking idea.

Additional Support for an Inhibitory Account of Retrieval-Induced Forgetting

Two other findings strongly implicate inhibitory processes in RIF. The first derives from the independent-probe technique introduced by Anderson and Spellman (1995). This technique was devised to test whether Rp– items are inhibited in the strong sense of the word—that is, whether access to those items would be impaired in general, not just in response to a category cue under which those items were studied in phase one of the experiment. Anderson and Spellman reasoned that if competing responses are truly suppressed during the retrieval attempt for a designated target, then these inhibited items should be more difficult to retrieve, not only from the studied retrieval cue, but also from other appropriate cues as well. Consistent with that conjecture, practice retrieving some exemplars of a studied category (e.g., *Red Blood*) impaired the later recall of unpracticed exemplars in that category (e.g., *Red Radish*) even when recall on the final test was cued by a novel category cue, such as *Food*—a cue that was not used in the retrieval practice phase.

A critical aspect of this finding is that the recall of the unpracticed exemplars in response to the unpracticed cue was impaired even though that cue was unrelated to the items strengthened during retrieval practice. The unpracticed cue, therefore, provides a measure of the accessibility of these related, unpracticed items that is independent of associative interference from the practiced targets. Using this same cue-independent technique but with very different materials—multidimensional geometric stimuli that could be categorized in terms of their location, shape, and color—Ciranni and Shimmamura (1999) also found evidenced supporting the inhibitory account of retrieval-induced forgetting. (See, however, Perfect et al., 2004, for a criticism of the logic behind the cue-independent technique.)

Finally, some newer findings also provide strong support for an inhibitory account. Storm, Bjork, Bjork, and Nestojko (2005) tested inhibitory versus noninhibitory accounts of RIF by introducing an interesting retrieval-practice condition—namely, one in which some cues, such as *Metals–Mu*________, posed an impossible retrieval task for the participants. That is, no studied member of the *Metals* category started with *Mu*, nor does any metal—at least among those known to the typical undergraduate—begin with those letters. Participants in Storm, Bjork et al.'s experiment, after having studied a typical Anderson et al. (1994) list of category-exemplar pairs, then received practice for half of the categories on the studied list. For half of these practiced categories, the retrieval practice involved retrieving new exemplars from that category, but ones that had not been presented in the studied list (e.g., *Fruit Or*________). For the other half of the practiced categories, the retrieval practice was impossible (e.g., *Metals Mu*________).

As expected, consistent with the findings of research by Bauml (2002), retrieving extralist members of a given category (the possible condition) resulted in RIF. Of more interest and greater theoretical importance, the impossible condition also produced RIF. Such a finding is especially difficult to explain via a noninhibitory account, such as blocking, because no member of the earlier-studied category is retrieved and strengthened in the impossible condition.

As Storm, Bjork et al. (2005) emphasize, it has often been assumed that retrieval practice must be successful in order for retrieval-induced forgetting to occur, but an inhibitory account does not require any such assumption. If, as is assumed in the inhibitory account, potentially interfering and competing information is suppressed in order to facilitate retrieval of target information, this process should occur whether the retrieval attempt succeeds or not. Indeed, aspects of Storm, Bjork et al.'s results suggest that the arguably more difficult impossible retrieval practice may have more RIF, consistent with the idea that a search for an impossible target might entail suppression of more competing responses than would the retrieval search for a possible target, which can be expected to cease as soon as a given desired target is retrieved.

Adaptive Consequences of the Retrieval Inhibition Assumed to Underlie RIF

As observed in the RIF paradigm, when information is retrieved from memory, it becomes more recallable than it would be otherwise, while other information associated with the same cue or configuration of cues becomes less recallable. As discussed earlier, these positive and negative effects of retrieval both illustrate the role of retrieval as a memory modifier (Bjork, 1994) and support the new theory of disuse (Bjork & Bjork, 1992) in which it is assumed that retrieval of an item associated to a given cue not only increases that item's retrieval strength but decreases the retrieval strength of other items associated with the same cue. We see these positive and negative effects of retrieval as being both adaptive and essential to an efficiently functioning memory system: Access to information likely to be needed again in the future is increased, while interference from competing information is reduced. Moreover, because it is the retrieval strength of the unpracticed information that is assumed to be inhibited or suppressed while its storage strength in memory remains unchanged—although retrieval access to such information may be temporarily impaired—should our circumstances change and we need to regain access to such information in the future, we can become fluent in its use again more quickly than were we to have to learn it anew from scratch. (For a more detailed discussion of the adaptive consequences of the retrieval inhibition or suppression assumed to underlie retrieval-induced forgetting—as well as some potentially negative consequences—the reader is referred to Bjork, Bjork, & MacLeod, 2006).

RESOLVING PART 1 AND PART 2: SOME SPECULATIONS

In the foregoing section, we argued that RIF is attributable to a selection/suppression mechanism and that RIF phenomena cannot be explained by blocking mechanisms of the type that emerged, 30 years or so ago, as the dominant explanation of phenomena such as output interference and the effects of part-list cueing. In this final section we consider the other side of the coin: What role might selection coupled with suppression play in output interference and part-list cueing, if any?

On the Possibility of Overlapping and Interacting Processes

As a kind of preamble to our speculations, it is important to emphasize that the various theoretical mechanisms that have been proposed to account for recall as a self-propagating or self-limiting process tend not to be mutually exclusive. Selection, suppression, spreading activation, blocking, sampling with replacement, associative chaining, covert mediation or rehearsal, and metacognitive processes, such as stopping rules, may all be involved, often in parallel, especially across the range of materials and procedures that have been used in laboratory studies.

Even in the case of research on RIF, the involvement of processes other than selection and suppression are clearly indicated. One such indication is the tendency for items weakly associated to a given cue to profit, rather than suffer, from retrieval practice on other items associated with that cue, which suggests that spreading activation, as well as selection and suppression, may be involved in the RIF paradigm. Another indication is the tendency for RIF to shift toward retrieval-induced facilitation as studied materials become more linked or integrated. Anderson and McCulloch (1999), for example, report multiple results on "integration as a boundary condition for retrieval-induced forgetting"—results that suggest that processes such as associative chaining during study and mediation during recall can result in unpracticed items being helped, not hindered, by retrieval practice on other items. Recently, and consistent with Anderson and McCulloch's findings, Chan, McDermott, and Roediger (in press) have found clear evidence of retrieval-induced facilitation when strong linkages exist between Rp+ and Rp– items. After having participants study an article on the biological characteristics and living habits of the toucan bird, Chan et al. structured the retrieval-practice phase so that the questions asked in that session (e.g., "Where do toucan birds sleep at night?") bore a relationship to other, unasked, questions (e.g., "What other bird species is the toucan related to?"). From the article it was clear that toucans are not themselves able to make holes in trees, but that they often sleep in holes made by woodpeckers, to which they are related. Given that linkage, final recall of the Rp– question, as well as the Rp+ question, was facilitated for these participants compared to that of control participants who only read the article and were not asked any questions about it prior to the final recall test.

Selection and Suppression in Output Interference

Extending the selection-plus-suppression idea to output-interference effects in free recall requires, first of all, the mostly uncontroversial assumption that all studied items are associated with a common list node or episodic context. The second necessary assumption is that recalling a given item from the list or study episode requires selecting that item from among all the items in that list or episode, which also seems uncontroversial. The final, and crucial, assumption is that the process of selection requires that competing items—that is, other list items that might be recalled—be suppressed. As free recall proceeds, then, the repeated suppression of yet-to-be-recalled items makes those items harder and harder to access and the recall process gradually grinds to a halt. When categorized lists are studied and recall is cued by category names the corresponding argument seems straightforward—and much the same as that advanced by Anderson et al. (1994) to explain the RIF observed with categorized materials: Each retrieval of a member of the category requires suppressing others and, gradually, makes the remaining members of the category hard to access.

Specifying the possible contributions of selection-plus-suppression to part-list cueing effects seems much less straightforward. To begin with, for example, RIF-type effects have not generally been found when opportunities to restudy some of the items presented earlier replace actual retrieval of those items during the retrieval-practice phase of the RIF paradigm (see, e.g., Anderson et al., 2000; Bauml, 2002; Ciranni & Shimamura, 1999). The restudied items benefit from additional study, but apparently not at a cost to the Rp– items. In terms of the selection-plus-suppression account of RIF, this lack of a cost to the Rp– items arises because such items do not compete for retrieval—and thus do not have to be selected against—when Rp+ items are re-presented as intact items for study as opposed to being retrieved.

In contrast, the typical part-list cueing procedure differs in possibly important ways from the RIF paradigm, particularly at the time of the final test, and these differences may provide opportunities—as outlined below—for selection-plus-suppression dynamics to play a role in the production of part-list cueing effects. (For a detailed comparison of procedural differences between the RIF, part-list cueing, and retroactive-interference paradigms, see Anderson & Bjork, 1994; Anderson et al., 1994). Typically, the effects of part-list cueing have been examined in the context of free recall, with participants told either to recall all list items, including the items presented as cues, or to recall only the remaining list items (i.e., targets), excluding the cue items. Given that the cue items have been re-presented, it seems safe to assume that they are the items most available for recall. In the case where they, too, are to be recalled, they tend to be recalled first, in which case their selection could be accompanied by suppression of the yet-to-be-recalled target items. When the cue items are not to be recalled, the argument becomes more complicated and other factors may well come into play, such as retrieval-strategy disruption (see Basden & Basden, 1995; Nickerson, 1984). The editing burden entailed by the instructions not to recall the strongest items (i.e., the items presented as cues) may, however, lead to a covert cycle of such items being recalled, but not written down, which then may introduce suppression dynamics such as those we have hypothesized may play a role in output interference. What can be said with some confidence is that inhibition owing to part-list cueing is not only highly unintuitive from a layperson's standpoint, but also an "enigma for memory researchers" (Nickerson, 1984)—in part, no doubt—because it is a product of multiple interacting dynamics.

CONCLUDING COMMENT

Our coverage of the earlier era of research on "recall as a self-limiting process" (Roediger, 1978) in the context of the current era of research on "retrieval-induced forgetting" (Anderson et al., 1994) illustrates that fundamental issues in memory research tend to reappear, clothed, often, in different paradigms. Our chapter also illustrates, as do the other chapters in the present volume, that—from his student days to the present—Henry (Roddy) Roediger III has been a mover and a shaker in virtually every successful effort our field has made to understand the complex, unintuitive, and multifaceted dynamics of human memory.

REFERENCES

Allen, G. A., Mahler, W. A., & Estes, W. K. (1969). Effects of recall tests on long-term retention of paired associates. *Journal of Verbal Learning and Verbal Behavior*, *8*, 463–470.

Anderson, J. R. (1972). FRAN: A simulation model of free recall. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 5, pp. 315–379). New York: Academic Press.

Anderson, M. C., Bjork, E. L., & Bjork, R. A. (2000). Retrieval-induced forgetting: Evidence for a recall-specific mechanism. *Psychonomic Bulletin and Review*, *7*, 522–530.

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). New York: 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., & McCulloch, K. C. (1999). Integration as a general boundary condition on retrieval-induced forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 608–629.

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.

Barnier, A. J., Hung, L., & Conway, M. A. (2004). Retrieval-induced forgetting of emotional and unemotional autobiographical memories. *Cognition and Emotion*, *18*, 457–477.

Basden, D. R., & Basden, B. H. (1995). Part-list cueing: A retrieval strategy disruption interpretation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 1656–1669.

Bauml, K.-H. (2002). Semantic generation can cause episodic forgetting. *Psychological Science*, *13*, 356–360.

Bjork, E. L., Bjork, R. A., & MacLeod, M. D. (2006). Types and consequences of forgetting: Intended and unintended. In L.-G. Nilsson & N. Ohta (Eds.), *Memory and society: Psychological perspectives* (pp. 134–158). Psychology Press: Hove, UK.

Bjork, R. A. (1975). Retrieval as a memory modifier: An interpretation of negative recency and related phenomenon. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola Symposium* (pp. 123–144). New York: Wiley

Bjork, R. A. (1994). Memory and metamemory considerations in the training of human beings. In J. Metcalfe & A. Shimamura (Eds.), *Metacognition: Knowing about knowing* (pp. 185–205). Cambridge, MA: MIT Press.

Bjork, R. A. (2001). Recency and recovery in human memory. In H. L. Roediger, J. S. Nairne, I. Neath, & A. M. Suprenant (Eds.), *The nature of remembering: Essays in honor of Robert G. Crowder* (pp. 211–232). Washington, DC: American Psychological Association Press.

Bjork, R. A., & Bjork, E. L. (1992). A new theory of disuse and an old theory of stimulus fluctuation. In A. Healy, S. Kosslyn, & R. Shiffrin (Eds.), *From learning processes to cognitive processes: Essays in honor of William K. Estes* (Vol. 2, pp. 35–67). Hillsdale, NJ: Lawrence Erlbaum Associates; Inc.

Chan, J. C. K., McDermott, K. B., & Roediger, H. L. (in press). Retrieval-induced facilitation: Initially nontested material can benefit from prior testing of related material. *Journal of Experimental Psychology: General*.

Ciranni, M. A., & Shimamura, A. P. (1999). Retrieval-induced forgetting in episodic memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 1403–1414.

Dong, T. (1972). Cued partial recall of categorized words. *Journal of Experimental Psychology*, *93*, 123–129.

Dunn, E. W., & Spellman, B. A. (2003). Forgetting by remembering: Stereotype inhibition through rehearsal of alternative aspects of identity. *Journal of Experimental Social Psychology*, *39*, 420–433.

Estes, W. K. (1955). Statistical theory of distributional phenomena in learning. *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. 161–190). Washington, DC: Winston.

Hull, C. L. (1943). *The principles of behavior*. New York: Appleton-Century-Crofts.

Landauer, T. K., & Bjork, R. A. (1978). Optimum rehearsal patterns and name learning. In M. M. Gruneberg, P. E. Morris, & R. N. Skykes (Eds.), *Practical aspects of memory* (pp. 625–632). London: Academic Press.

Macrae, C. N., & MacLeod, M. D. (1999). On recollection lost: When practice makes imperfect. *Journal of Personality and Social Psychology*, *77*, 463–473.

Mandler, G. (1967). Organization and memory. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation* (pp. 327–372). New York: Academic Press.

Melton, A. W. (1963). Implications of short-term memory for a general theory of memory. *Journal of Verbal Learning and Verbal Behavior*, *2*, 1–21.

Nickerson, R. S. (1984). Retrieval inhibition from part-set cueing: A persisting enigma in memory research. *Memory and Cognition*, *12*, 531–552.

Perfect, T. J., Stark, L., Tree, J. J., Moulin, C. J. A., Ahmed, L., & Hutter, R. (2004). Transfer appropriate forgetting: The cue-dependent nature of retrieval-induced forgetting. *Journal of Memory and Language*, *51*, 399–417.

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 and Cognition*, *2*, 261–269.

Roediger, H. L. (1978). Recall as a self-limiting process. *Memory and Cognition*, *6*, 54–63.

Roediger, H. L., & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science*, *17*, 249–255.

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

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

Saunders, J., & MacLeod, M. D. (2002). New evidence on the suggestibility of memory: The role of retrieval-induced forgetting in misinformation effects. *Journal of Experimental Psychology: Applied*, *8*, 127–142.

Shaw, J. S., III, Bjork, R. A., & Handal, A. (1995). Retrieval-induced forgetting in an eyewitness-memory paradigm. *Psychonomic Bulletin and Review*, *2*, 249–253.

Shiffrin, R. M. (1970). Memory search. In D. A. Norman (Ed.), *Models of memory* (pp. 375–447). New York: Academic Press.

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

Slamecka, N. J. (1969). Testing for associative storage in multitrial free recall. *Journal of Experimental Psychology*, *81*, 557–560.

Slamecka, N. J. (1984). Commentary on "An examination of trace storage in free recall." *Citation Classics*, *25*, 96. (Original work published in *Journal of Experimental Psychology*, *76*, 504–513).

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

Smith, A. D., D'Agostino, P. R., & Reid, L. S. (1970). Output interference in long-term memory. *Canadian Journal of Psychology*, *24*, 85–89.

Storm, B. C., Bjork, E. L., & Bjork, R. A. (2005). Social metacognitive judgments: The role of retrieval-induced forgetting in person memory and impressions. *Journal of Memory and Language*, *52*, 535–550.

Storm, B. C., Bjork, E. L., Bjork, R. A., & Nestojka, J. F. (in press). Is retrieval success a necessary condition for retrieval-induced forgetting? *Psychonomic Bulletin and Review*.

Thorndike, E. L. (1914). *The psychology of learning*. New York: Teachers College.

Tulving, E. (1966). Subjective organization and the effects of repetition in multitrial free recall learning. *Journal of Verbal Learning and Verbal Behavior*, *5*, 193–197.

Tulving, E., & Arbuckle, T. Y. (1963). Sources of intratrial interference in immediate recall of paired associates. *Journal of Verbal Learning and Verbal Behavior*, *1*, 321–334.

Tulving, E., & Osler, S. (1968). Effectiveness of retrieval cues in memory for words. *Journal of Experimental Psychology*, *77*, 593–601.

Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, *5*, 381–391.

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

Whitten, W. B., & Bjork, R. A. (1977) Learning from tests: The effects of spacing. *Journal of Verbal Learning and Verbal Behavior*, *16*, 465–478.