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: Erlbaum.

2 A New Theory of Disuse and an Old Theory of Stimulus Fluctuation

Robert A. Bjork
Elizabeth Ligon Bjork
University of California, Los Angeles

Speakers at the William K. Estes Symposium at Harvard University were asked to pick, if possible, a research topic where they could trace the influence of W. K. Estes in the work to be reported at the symposium. In the first author's case, that did not narrow down the possible topics in any substantial way. The work that seemed most timely to report at the symposium, however—the collaborative effort we refer to herein as a "new theory of disuse"—seemed not to be a particularly good example of the various significant influences William K. Estes has had on the two of us. Upon reflection, however, certain formal aspects of our theory correspond to a version of Estes' stimulus sampling theory, a version that incorporates what we consider to be one of the great insights in the history of research on learning and memory. That insight, initially reported in two short papers in the 1955 volume of the *Psychological Review* (Estes, 1955a, 1955b), is implemented in the so-called stimulus fluctuation version of Estes' statistical theory of learning.

However delayed and unconscious the influences may have been, we feel that our new theory of disuse owes some of its features to Estes' theory of stimulus fluctuation. One goal of this chapter is to sketch the similarities and differences between our theory-of-disuse framework and Estes' fluctuation model. In the sections that follow, we first summarize the characteristics of human memory that we feel suggest the storage and retrieval properties we postulate in our theory of disuse. We then present that framework along with some of its predictions and some arguments why such a pattern of storage and retrieval characteristics might be, overall, adaptive. We conclude with a section in which we first describe and pay homage to Estes' stimulus fluctuation insight, and we then compare and contrast the fluctuation model and our theory of disuse.

SOME IMPORTANT PECULIARITIES OF STORAGE AND RETRIEVAL PROCESSES IN HUMANS AND ANIMALS

In some important respects, the storage and retrieval properties of the human memory system differ markedly from the storage and retrieval properties of man-made memory devices, such as a tape recorder or the memory in a computer. The particular properties of human memory we summarize in the following sections have some important implications, in our view, as to the overall functional organization of the memory system.

Our Differing Capacities for Storage and Retrieval

In the modal view of human long-term memory, new information is stored in memory—not by recording some literal copy of that information, but, rather, by interpreting that information in terms of what we already know. New items of information are "fit in" to memory, so to speak, in terms of their meaning; that is, in terms of their semantic relationship to items, schemas, and scripts already in long-term memory. The process appears to be one of virtually unlimited capacity. Rather than using up storage capacity, as would be the case if our long-term memories were some kind of box or tape, the act of storing new information in memory appears to create the opportunities for additional storage: The more knowledge we have in a given domain, the more ways there are to store additional information. It also appears, at least to a first approximation, that once "entrenched" in long-term memory, information remains in memory for an indefinitely long period of time. (See Landauer, 1986, for some impressive estimates of the functional capacities of human long-term memory in some separate domains of knowledge; see, also, Miller & Gildea, 1987, for an impressive example of knowledge accumulation in a given domain—vocabulary items—across the first 18 years or so of life.)

As impressive as we appear to be in terms of getting information into long-term memory, we are far less impressive at getting information out of the system. The retrieval process is erratic, highly fallible, and heavily cue dependent. Information (such as a name, phone number, or street address) that is recallable on one occasion without apparent effort can be impossible to recall on another occasion. Even the most highly-overlearned and frequently accessed information, such as a street address or phone number, eventually becomes nonretrievable over years of disuse, though it is a simple matter via tests of relearning or recognition to demonstrate that such items still exist in storage.

What we can and cannot recall at a given point in time appears to be governed by the cues that are available to us, where such "cues" may be environmental, interpersonal, emotional, or physical (body states) as well as ones that bear a direct associative relationship to the target item. Cues that were originally associated in storage with the target information need to be reinstated, physically or mentally, at the time of retrieval. The importance of such reinstatement is, of course, at the heart of Tulving and Thompson's (1973) encoding specificity principle. (For a discussion of the importance of mental rather than physical reinstatement, see Bjork and Richardson-Klavehn, 1989.)

Retrieval as a Memory Modifier

In contrast to a computer or a tape recorder, where retrieving stored information leaves the stored representation of that information—and other information—unchanged (except for some slight physical deterioration, if any, that may result), the act of retrieving information from human memory modifies the system. The information retrieved becomes more retrievable in the future, and other information can become less retrievable.

In terms of its positive consequences, the act of retrieval is itself a potent learning event. As an overall generalization, the act of retrieving an item of information is considerably more potent in terms of facilitating its subsequent successful recall than is an additional study trial on that item. The actual extent to which a *successful* retrieval facilitates later retrieval appears to depend on how difficult or involved that act of retrieval is, with the subsequent benefits being a positive function of retrieval difficulty.

The foregoing assertions about the positive consequences of retrieval are based on a huge body of empirical research dating back decades (an abbreviated list of references is Allen, Mahler, & Estes, 1969; Bartlett, 1977; Bjork, 1975; Gardiner, Craik, & Bleasdale, 1973; Hogan & Kintsch, 1971; Izawa, 1970; Landauer & Bjork, 1978; Modigliani, 1976; Thompson, Wenger, & Bartling, 1978; and Whitten & Bjork, 1977). The assertion that there may be negative consequences for other items in memory, however, is a more novel idea that requires some more explicit evidence. Experiments by Bjork and Geiselman (1978) and by Richardson-Klavehn (1988) have the nice property of demonstrating simultaneously the positive (retrieval practice) and negative (retrieval competition) effects of retrieval.

In the Bjork and Geiselman experiment (1978, Experiment 3), subjects were presented two lists, each of which consisted of 12 word pairs (i.e., there were 24 words in each list). Within each list, after a filled delay of 6 s or 12 s, subjects were cued, in unpredictable fashion, whether to continue to remember or to forget each word pair. The consequences of the forget/remember and delay manipulations were of primary interest to Bjork and Geiselman, but are not relevant to the present issue. What *is* relevant is that each list was followed either by an immediate test of free recall for the to-be-remembered pairs or by a comparable period of deductive reasoning problems, and there was a final (end of experiment) test of free recall for all the words in both lists.

In Fig. 2.1 the final free recall of List 1 (bottom pair of curves) and List 2 (top

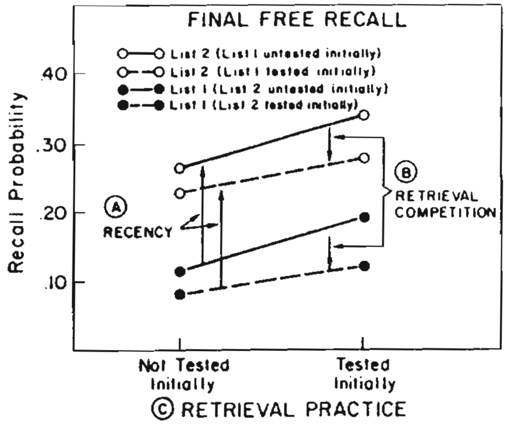

FIG. 2.1. Final free-recall probabilities as a function of list position and initial-test conditions. The influence of list recency (A), retrieval competition (B), and retrieval practice (C) is indicated. (After Bjork & Geiselman, 1978).

pair of curves) are shown as a function of the four possible combinations of whether List 1 and List 2 were tested or not tested. There was a recency effect (List 2 recalled better than List 1), but independent of that effect, each list was better recalled if it had been tested initially (the right hand points vs. the left hand points) and if the other list had not been tested initially (the top curve for each list vs. the bottom curve for each list). Thus, an initial free recall test on a given list not only facilitated its final free recall, but also impaired the final free recall of the other list.

Richardson-Klavehn (1988, Experiment 3) obtained analogous results with a different paradigm. Subjects were presented a list of 40 word-word paired associates to study. The stimulus and response words in each pair were low associates of each other, and the list of 40 pairs was presented three times in succession. At the end of the study phase, to give subjects a sense of closure, 10 of the 40 pairs were tested by presenting a given stimulus word as a cue for its associated response word. After a 5-min break, subjects were asked to learn a second list of 40 paired associates that consisted of the same 40 stimulus words paired with new response words, which were also low associates of the stimulus words (the lists, therefore, bore an AB-AC relationship to each other). List 2 was also presented three times for study and at the end of the study phase, 10 stimuli (a different 10 stimuli) were presented as a cued-recall test for their corresponding response words. Finally, 24 hours later, subjects' cued recall of List 1 response words was tested by presenting all 40 stimulus words one at a time.

Richardson-Klavehn was actually interested in the degree to which the retroactive interference owing to List 2 in the recall of List 1 items was a function of whether List 2 was learned in the same physical setting as List 1. His design, however, permits one to look at the 24-hr delayed recall of List 1 items for which neither the List 1 nor List 2 response was tested on Day 1 (20 pairs) versus the

List 1 items for which only the List 1 or the List 2 response was tested on Day 1 (10 pairs in each case). Collapsed across the same environmental manipulation, 63% of the List 1 responses were recalled when neither response had been tested on Day 1. When the List 1 response had been tested on Day 1, performance rose to 84%, and—of more interest—when the List 2 response had been tested on Day 1, performance *dropped* to 49%.

In both the Bjork and Geiselman and the Richardson-Klavehn experiments, then, an initial act of retrieval facilitated later retrieval of those items while impairing the recall of other items. In the former case, that pattern appeared on a whole-list basis in delayed free recall where the only obvious relationship between the lists was their common episodic context. In the latter case, that pattern appeared on a pair-specific basis where the List 1 and List 2 response words each bore a weak semantic relationship to a given stimulus word.

Relative or Absolute Accessibility to Items in Memory Regresses over Time

A final important peculiarity of human memory is that, over time, the accessibility of memory representations constructed earlier tends to increase relative to the accessibility of related memory representations constructed later. That is, a kind of regression process appears to be a fundamental factor in determining which items are and are not accessible (retrievable) from storage. When we update our memory representations (by learning to operate a new car, for example, or by learning a new golf swing, or a new word processing program, or the new married name of a friend, or a new list in a memory experiment), it is the new representation that is most accessible at the end of that learning process. With disuse of both representations, however, the pattern changes: There is loss of access to the more recent representation and a recovery in access (sometimes in absolute as well as relative terms) to the earlier representation.

The evidence for such memory regression in the real world is mostly observational and anecdotal. Coaches and skilled athletes tend to be aware that a long layoff leads to recovery of old habits. A layoff can even help a skilled athlete recover from a recent problem, whereas a layoff can lead to a major step backwards for an amateur who has been improving rapidly (the regression would be to good and not-so-good habits, respectively, in those two cases). People in the military tend also to be aware, in a general way, that regression is a problem. Personnel trained in new procedures and equipment can appear to be well-trained but nonetheless take inappropriate actions (that is, actions appropriate to the old equipment or procedures) at a later time, particularly under stress, as in the heat of battle.

Such regression processes no doubt generalize beyond the military and sports worlds. Bjork (1978) has argued that the common experience of being surprised "at how much a child has grown up, a friend has aged, or a town has changed since the last time one saw that child, friend, or town" (p. 250) may be interpretable in terms of the regression of memory representations.

> Children do grow up, of course, friends do age, and towns do change, but a subjective judgment of such changes based on the difference between a regressed memory representation and the current state of the child, friend, or town will overestimate the actual changes. One particularly compelling instance of such overestimation, in my view, occurs when one is away from one's small children for a week or two. The apparent growth can far exceed any actual growth that could have taken place in that time. The fact that a day or so later such phenomenal growth is no longer apparent argues against the reality of such apparent changes. (p. 250)

Beyond such real-world observations and anecdotes, there exists solid experimental evidence for such a regression process in humans and animals. In classic list-learning paradigms in the interference tradition (such as the AB-AC paired-associate paradigm), the pattern of proactive and retoractive effects changes markedly with the retention interval following List 2 learning. With the AB-AC paradigm, for example, retroactive interference (impaired recall of AB pairs owing to AC learning) tends to decrease with retention interval, whereas proactive interference (impaired recall of AC pairs owing to AB learning) tends to increase with retention interval (see, for example, Briggs, 1954; Ceraso & Henderson, 1965, 1966; Forrester, 1970; Koppenaal, 1963; Postman, Stark, & Fraser. 1968). In a number of those experiments employing MFR (give the first response that comes to mind) or MMFR (give all responses that you can remember) tests, there was an absolute increase in List 1 responses with an increasing retention interval following List 2 learning. Such "spontaneous recovery" is, of course, a common finding with extinction paradigms in animal research.

Another type of result that we feel demonstrates memory regression is the change from recency to primacy with delay. If subjects are asked—at the end of an experimental session during which a series of lists has been presented—to recall all the items they can remember from any of the lists, there is typically a strong list-recency effect: More items are recalled from the last list than from the next-to-last list, and so forth. After 24 hours (Bjork & Whitten, 1974) or a week (Bower & Reitman, 1972), however, if subjects are again asked to recall the lists, their recall exhibits primacy: The first list is recalled better than the second list, and so forth. Such a recency-to-primacy change with delay appears to occur across a broad range of time scales, types of tasks and materials, and species.

A particularly dramatic cross-species demonstration on a short time scale is provided by the work of Wright and his co-workers (Santiago & Wright, 1984; Wright, Santiago, & Sands, 1984; Wright, Santiago, Sands, Kendrick, & Cook, 1985; for a discussion see Wright, 1989). Pigeons, monkeys, and humans were given analogous memory-search tasks with visual materials. On each trial, four stimuli were shown one at a time (pictures in the experiments with pigeons and monkeys, kaleidoscope patterns in the experiments with humans). A probe item was then presented at various delays following the last item in a given memory set; the organism was required to make one response if the probe item had been in the memory set and another response if it was a new item. For "old" items (i.e., probes that matched a memory-set item), a plot of percent correct responding as a function of input serial position and probe delay exhibits a strikingly similar pattern for the pigeons, monkeys, and humans. At the shortest (zero s) delay, there is a monotonic recency effect for all three species; at the longest delay (10 s for pigeons, 30 s for monkeys, 100 s for humans), there is a monotonic primacy effect for all three species; and at intermediate delays, the functions are U-shaped for all three species. In addition to that change of pattern, there is an *absolute* increase in performance on the first memory-set item from the shortest to the longest probe delay.

A NEW THEORY OF DISUSE

In this section we sketch the assumptions, predictions, and adaptive features of a new theory of disuse. The assumptions of the theory emerge, in large part, from our analysis of the "peculiarities" of human memory summarized in the preceding section. We give those phenomena heavy weight because we feel (a) that they reflect an architecture of the memory system that differs from the architecture of man-made memory systems, and (b) that they imply storage and retrieval properties that are adaptive in the overall day-by-day functioning of human memory. In the latter respect, our approach corresponds in a general way to the approach advocated by Anderson (1989), who argues for "an alternative way of casting a theory of memory" based on his *Principle of Rationality* ("Human memory behaves as an optimal solution to the information-retrieval problems facing humans" [p. 195]). Baddeley (1988), too, has asked researchers to keep in mind that ". . . man is a biological and social animal, and like other animals is the product of evolution" (p. 15).

In formulating the theory, we take as a starting point the fact that items of information, no matter how accessible and overlearned they may be at some point in time, eventually become nonrecallable with disuse. In a symposium on autobiographical memory at a conference on practical aspects of memory held in Wales in the summer of 1987, we (Bjork & Bjork, 1988) argued that it is a mistake to view the frequent retrieval failures that characterize human memory as simply a weakness of the system. In terms of keeping the system current—in part *because* our storage capacity is so enormous—it is adaptive that we lose retrieval access to information with disuse. In attempting to recall our current home address, it is not useful to recall from our memories every home address we may have had in the past. That those prior addresses have, for the most part, become nonretrievable (though they remain stored) reduces confusion and speeds access to our current address.

A variety of such considerations led us to formulate a new theory of disuse. Thorndike's (1914) original law of disuse, of course, stands as one of the most thoroughly discredited of the various "laws" psychologists have put forward over the years—which is a considerable distinction. Thorndike argued that learned habits, without continued practice, fade or decay from memory with the passage of time. Interference theorists (e.g., McGeoch, 1932) presented compelling evidence against the notion that memory traces decay and that such decay takes place as a function of time alone. We, however, believe—the fate of the original law of disuse notwithstanding—that the "important peculiarities" of human memory that we have summarized suggest a modified theory of disuse. The following assumptions represent an attempt to specify in considerably more detail than we did initially the assumptions of such a theory.

Assumptions of the Theory

1. An item in memory can be characterized by two "strengths," a storage strength and a retrieval strength. The former measures, in a general way, how well learned an item is; the latter measures the current ease of access to the item in memory. The probability that an item can be recalled in response to a given cue is completely determined by its retrieval strength (and on the retrieval strength of other items in the set of items associated to that cue, as specified in Assumption 3) and is independent of its storage strength. That is, storage strength is a latent variable that has no direct effects on performance. Items with high storage strength can have low retrieval strength (e.g., a home phone number one had for 5 years 20 years ago), and items with low storage strength can have high retrieval strength (e.g., one's room number at a resort hotel on, say, the third day of one's stay at that hotel).

(Viewed historically, the distinction we make between storage strength and retrieval strength is analogous to distinctions made by learning theorists of another era. Hull [1943] distinguished between habit strength and momentary reaction potential; Estes [1955b], as we discuss in more detail later, distinguished between habit strength and response strength; even Skinner [1938], distinguished between reflex reserve and reflex strength.)

2. The storage strength of a given item grows as a pure accumulation process, that is, as a monotonic function of opportunities to study or recall that item. To at least a first approximation, therefore, it is assumed that storage strength once accumulated is never lost. One consequence of this assumption is that there is no limit on the amount of information that can be stored in long-term memory; that is, on the sum of storage strengths across items. The growth function for a given item is, however, assumed to be negatively accelerated: The increment in storage strength for a given item owing to a study or test event is assumed to be a decreasing function of its current storage strength. We also assume that the increment in storage strength is a decreasing function of its current retrieval strength; that is, high retrieval strength is assumed to retard the accumulation of additional storage strength.

3. Whereas there is no limit on storage capacity, there is a limit on retrieval capacity; that is, on the total number of items that are retrievable at any one point in time in response to a retrieval cue or set of cues. In fact, we assume two kinds of limits on retrieval capacity. The first, an overall limit, simply reflects our assertion (Assumption 5) that retrieval strength, in contrast to storage strength, is lost as a function of subsequent study and test events on other items. Thus, at some point, a kind of dynamic equilibrium is reached where any gain in retrieval strength for items being studied or tested is offset by a corresponding loss in retrieval strength summed across other items in memory.

The second, and often more significant limitation, arises from how we characterize the cue-dependent nature of retrieval. For an item to be recalled in response to a given cue, we assume that its representation must be discriminated from the representations of other items in memory associated to that cue, and that it must be reconstructed or integrated from its representation. Discriminating an item is assumed to be a function of its retrieval strength relative to the strengths of other items in the cued set (its strength, e.g., normalized with respect to the sum of the item strengths in the cued set). Reconstructing the item for output, on the other hand, is presumed to be a straightforward function of its absolute retrieval strength. The net effect of these assumptions is to put a limit on the number of items that can be accessible in memory at a given time. As items are added to memory, or as the retrieval strengths of certain items are increased, other items become less recallable.

4. The act of retrieving an item from memory, and the act of studying an item, both result in increments to that item's retrieval strength as well as to its storage strength, but retrieval is the more potent event. That is, the act of retrieving an item (successfully) results in larger increments to storage strength and retrieval strength than does studying an item. Increments in retrieval strength, in either case, are assumed to be a decreasing function of an item's current retrieval strength and an increasing function of its current storage strength. One consequence of these assumptions is that the benefits of a successful retrieval, in terms of its influence on that item's subsequent retrieval strength, are larger the more difficult or involved the act of retrieval (low retrieval strength) and the better registered the item is in memory (high storage strength).

5. The decrement in an item's retrieval strength owing to the learning or retrieval of other items is, conversely, assumed to be greater the higher the item's current retrieval strength and the lower the item's current storage strength. An-

other way to characterize these assumptions is to say that the gain and loss of retrieval strength are both negatively accelerated, and that storage strength acts to enhance the gain and retard the loss of retrieval strength.

It follows from the foregoing assumptions that increasing the retrieval strengths of certain items (via study or test events) makes other items less retrievable. Such competitive effects will tend to be governed by similarity or category relationships defined semantically or episodically. That is, a given retrieval cue or combination of cues will define a set of associated items in memory—which, in general, will bear some type of similarity relationship to each other—and the dynamics of competition for retrieval capacity, as was spelled out in Assumption 3, take place across that set.

"Predictions" of the Theory

Although we stay at the qualitative level in this chapter, it is, of course, our goal to predict the kind of phenomena we summarize below quantitatively as well as qualitatively. As may be apparent from how the assumptions of our theory are stated, one straightforward, though surprisingly complex, mathematical implementation of the theory is in terms of a set of interdependent linear difference equations. That representation, judging from our preliminary efforts to derive certain predictions, appears quite promising.

The "important peculiarities" of human memory discussed earlier influenced heavily the assumptions of our new theory of disuse, so it is not surprising or impressive that the theory can accommodate those findings. Certain aspects of the theory's account of those and other properties of human memory, however, warrant comment.

Storage versus Retrieval. The heroic storage capabilities and relatively frail retrieval capabilities of human memory emerge directly from the assumptions. The cue-dependent nature of retrieval processes is built into the theory as well. It should also be apparent that an item in memory, however well learned (that is, however high its storage strength), becomes inaccessible if not retrieved periodically. As other items are learned or strengthened via retrieval practice, the item in question becomes less and less accessible. It is not staying competitive, so to speak, for the limited retrieval capacities that characterize the system. The gradual (or, under some circumstances, not so gradual) loss of retrieval access is not a consequence of the mere passage of time, but, rather, is a consequence of the learning and practice of other items. It is worth noting that the memory-as-a-box analogy, which is so wrong on the storage side (as argued earlier) may have some truth on the retrieval side. As we make some items in memory more and more accessible, according to our theory of disuse, there is less and less remaining retrieval capacity for other items. This viewpoint, then, may exonerate the

ichthyologist David Starr Jordon, who as President of Stanford University complained that every time he learned the name of a student, he forgot the name of a fish. He is often cited uncharitably as someone who had a fallacious idea of the capacity of human memory. Given the limit on retrieval capacity assumed in our theory, however, an ichthyologist suddenly spending considerable time learning and retrieving the names of a large number of students could well lose access to the names of certain fish.

Retrieval Practice and Retrieval Competition. That retrieving an item from memory facilitates the later retrieval of that item, and that such effects are greater the greater the difficulty of the initial retrieval, is simply assumed in the theory. The competitive effects of retrieval—that is, that certain other items in memory become *less* retrievable in the future—arise because of the cue-specific and global limitations on retrieval capacity. As some items in the set of items associated to a given cue gain in retrieval strength, other items in that set will suffer, in part because they become less discriminable or "findable" in memory owing to their lessened *relative* retrieval strengths, and in part because their *absolute* retrieval strength will decrease with their disuse and the "use" of the other items in the set.

Regression/Recovery Phenomena. The relative or absolute recovery of older memory representations with time emerges from the theory in a somewhat less obvious way. It may be easiest to explicate the basic mechanisms in the context of classic unlearning and spontaneous-recovery phenomena (e.g., Melton & Irwin, 1940; Postman et al., 1968). The unlearning of a first list as a consequence of learning a similar second list (say, e.g., in an AB-AC paired-associate paradigm) comes about in the theory because increasing the accessibility of second-list responses decreases the accessibility of first-list responses. With the passage of time, however, and the disuse of both List 1 and List 2, the loss of List 1 retrieval strength will be offset by a greater loss of List 2 retrieval strength (because the loss of retrieval strength is negatively accelerated).

Whether there is an absolute increase in the recall of List 1 responses as the retention interval increases, or only an increase relative to the recall of List 2 responses, depends in the theory on certain parameter values. In particular, if List 1 responses can be viewed as having achieved a higher level of storage strength than List 2 responses (which could happen, e.g., if List 1 responses are rehearsed during List 2 learning), then the recovery of List 1 responses could be substantial in absolute terms. In that case, as asserted in Assumption 5, the higher storage strength of List 1 responses will slow the rate of loss of their retrieval strengths; at some value of the delay since List 2 learning, the retrieval-strength curves for List 1 and List 2 responses will cross, with List 1 responses having higher absolute retrieval strength after that point.

The change from recency to primacy as a function of retention interval is predicted for similar reasons. Once again, the absolute increase in performance on initial lists or items with delay (as in the pigeon/monkey/human data discussed by Wright, 1989) is explicable in the theory by assuming that early items, via cumulative rehearsal or other factors, reach a higher level of storage strength than do later items. The kind of anecdotal regression phenomena cited earlier for skilled athletes and military personnel may constitute cases where it is particularly plausible to assume that early habits have a higher level of storage strength in memory than do later habits. One's new, improved tennis serve is not going to be as entrenched in memory as one's old tennis serve, for example; nor is it likely that the operator of a new piece of military hardware has had the same amount of experience with its controls as he or she had with the controls of the hardware it replaced.

Overlearning and Repeated Learning. It is a time-honored result in both the human and animal literature that additional learning trials given after perfect performance is achieved (overlearning), or additional relearning sessions where performance is brought back to the original criterion (repeated learning) act to slow the rate of subsequent forgetting (e.g., Ebbinghaus, 1885/1964; Krueger, 1929). In the present theory, such effects have a straightforward interpretation in terms of the distinction between storage strength and retrieval strength. Performance is a function of retrieval strength, and performance cannot go beyond 100%. Storage strength can continue to accrue, however, as a function of overlearning or repeated learning, and increased storage strength acts (in the theory) to slow the loss of retrieval strength (and, hence, the observed forgetting).

This type of explanation of the effects of overlearning and repeated learning is not new or unique to our theory of disuse. In fact, Woodworth and Schlosberg's (1954) interpretation of the effects of repeated learning, as quoted below, corresponds quite closely to our interpretation (if one thinks of "trace strength" as storage strength and "readiness" as retrieval strength).

> On each successive day [the subject] learns to the same standard of one correct recitation. At the end of each day, he has reached the same degree of mastery. Why then should not forgetting proceed at the same rate? We are forced to conclude that the trace becomes stronger and stronger with each relearning. What is the same at the end of each day's learning is not the trace but the immediate recitability or recallability of the lesson, and recall obviously depends not alone upon the trace but also upon the momentary condition of readiness. Readiness depends very much on recency of impression. (p. 730)

Distribution of Practice. Our theory of disuse can account for the benefits on long-term retention of spacing repeated study trials, and for the interaction of repetition spacing and retention interval. At a formal level, the explanation is much the same as that provided by Estes' fluctuation model, which is outlined later in this chapter. In general, spacing of repetitions results in higher storage

strength than does massing of repetitions, which in turn slows the rate of loss of retrieval strength and, therefore, enhances long-term performance. Massing, however, can produce a higher initial level of retrieval strength, which, given a short enough retention interval, can result in a higher level of recall than that produced by spaced repetitions.

It is worth mentioning that the theory of disuse gives a natural interpretation of Jost's (1897) Law: If two associations are of equal strength but different ages, further study has greater value for the older one. In general, for two associations differing in age to be of equal retrieval strength, the older association must be at a higher level of storage strength. In that case, the increment in retrieval strength will be larger for the older association.

A Test of the Cue-Specific Nature of Retrieval Competition

The primary locus of retrieval competition, as specified earlier in Assumption 3, is within the set of items defined by a retrieval cue or combination of cues. Access to weaker items in the set is inhibited or obscured by the presence of stronger items in the set (see, e.g., the list strength experiments of Ratcliff, Clark, & Shiffrin, 1990). To ascertain whether such inhibitory or competitive effects are indeed cue (or category) specific, Anderson and Bjork (1990) have carried out several experiments using a new paradigm devised by Michael Anderson at UCLA. The basic procedure is as follows. Subjects are presented a list of pairs to study, each of which consists of a category name and an instance of the category (e.g., TREE Maple or DRINKS Scotch). Typically, six instances of eight different categories are presented in unblocked fashion producing a total of 48 pairs. There is then a retrieval-practice phase during which half the exemplars of half the categories (i.e., three members of each of four categories) are each tested three times. The retrieval of a given category member is induced by the category name and letter cues (e.g., TREE Ma_____), and the several retrievals of each of the 12 practiced items are interleaved with the tests on other items in such a fashion as to produce an expanding sequence of intertest intervals for each item (which appears to maximize the consequences of retrieval practice; see Landauer & Bjork, 1978). After an additional delay, subjects are given a final test in which they are cued with each category name and asked to free recall as many members of that category as they can remember having been presented in the original study list.

From the standpoint of the final test, a given category may have been accessed during the intervening retrieval-practice phase (an RP category) or it may not have been accessed (an NRP category). If it was a practiced category, then items within the category may have been practiced (RP+ items) or not practiced (RP− items).

This paradigm bears some relationship to part-list cuing procedures. At issue is the level of recall for RP+, RP−, and NRP items. Our theory predicts not only that RP+ items will be better recalled than NRP (and RP−) items, but also that NRP items will be recalled *better* than RP− items. The several intervening retrievals of RP+ items during the retrieval-practice phase will not only enhance their retrieval strength, but should also create problems for RP− items. The category name coupled with the instruction to recall members of that category presented during the study episode constitute a combined cue that designates a set of items in memory. For an NRP category, those items are on equal footing, so to speak, but for an RP category the RP− items are put at a competitive disadvantage by virtue of the presence of RP+ items in the cued set.

It is worth noting that the key prediction of the theory, that RP− items will be recalled more poorly than NRP items, runs counter to the effects of any episodic or semantic spread of activation to RP− items during the retrieval-practice phase. That is, the process of responding to TREE Ma_____ should, in semantic network models, spread some activation to other traces, particularly to other trees that were on the list and also bear some semantic relationship to *maple* (e.g., *oak*).

The basic pattern of results obtained by Anderson and Bjork (1990) is consistent with the theory of disuse. The size of the retrieval-practice effect (RP+ items vs. NRP items) and the retrieval inhibition effect (NRP items vs. RP− items) have varied as a function of category-to-exemplar associative strength and exemplar-to-category associative strength, but averaging across such manipulations the overall proportion of RP+, NRP, and RP− items recalled is .73, .49, and .37, respectively. Thus, the primary negative impact of the retrievals of RP+ items is on RP− items in the same category, not on NRP items in another category (although they, too, were in the original study list with RP+ items). To be able to say whether there are *any* inhibitory effects on NRP items requires a control condition (no retrieval practice phase at all) that has yet to be tested.

Adaptive Consequences for Autobiographical/Real-World Memory

We stated earlier our belief that the storage and retrieval properties of human memory define a system that is, all things considered, adaptive. In terms of the demands the world places on our memories, there is much to be said for a system with the properties embedded in our new theory of disuse. In general, the theory says that the items in memory that are most readily accessible to us are those we have been using (retrieving) lately. The items that have been retrieved frequently in the recent past will tend to be items highly relevant to our current interests, problems, goals, and general station in life. Statistically, those items will be relevant to the near future as well.

What about the items (names, numbers, facts, . . .) that we stop using? The theory says that eventually, however high their storage strength, they become nonretrievable. But why would we stop, for example, using the maiden name of a female friend or the combination on our gym locker? There are a variety of reasons, but they break into two types: Either that name or number is replaced (our female friend assumes a new married name; we get a new locker with a different combination) or events conspire to make us lose contact with our friend or locker (we move, she moves; we abandon hope and stop using the gym, and so forth). In either case, there are adaptive consequences of our losing retrieval access to that name or number. In the replacement case, the reason is obvious: We do not want to intrude our friend's maiden name or our old combination when trying to recall the current name or number. Items that are nonretrievable are also noninterfering (we do not, for example, typically have much difficulty recalling our current home phone number, even though we may have had many phone numbers in the past, some for a number of years, and even though we could verify that those out-of-date numbers still reside in our memory).

The adaptive aspects of losing retrieval access to names and numbers that we stop using by virtue of moving or changing our lives in some other way are a bit less obvious. If we consider, however, that for reasons of retrieval speed as well as accuracy we do not *want* everything in memory to be accessible, the adaptive consequences are fairly clear. The changes that result in our not using certain names and numbers will tend to require that we start using new and different names and numbers. Consider a new graduate student moving across the country to begin his or her graduate work. A whole panoply of new names (of people, streets, restaurants, campus buildings, . . .) and numbers (phone numbers, addresses, office numbers, ATM numbers, course numbers, . . .) must be learned and used—in addition to the learning that is a formal part of that student's graduate education. Those new names and numbers will gradually gain in storage strength and retrieval strength, and corresponding names and numbers from that student's past life will gradually lose retrieval strength. That pattern, overall, is adaptive. Put another way, *because* the storage capacity of human memory is so enormous, we need something to filter or direct access to the subset of items most relevant to our current needs. In our theory of disuse, retrieval strength plays that role.

It is also adaptive that items lose retrieval strength rather than storage strength with disuse: Because they become nonretrievable, they are noninterfering, but they are relearnable at an accelerated rate should they become pertinent again. (In the theory, the largest increments in retrieval strength as a consequence of studying or recalling an item take place for items that are high in storage strength and low in retrieval strength.) Such items may be recognizable as well. We have yet to characterize the process of recognition in our theory, but, presumably, recognition is sensitive to storage strength and is less dependent on retrieval strength than is recall. As an updating mechanism, then, the loss of retrieval strength for out-of-date items has several advantages over decay or displacement processes that leave no representation of the item in memory (such as the kind of destructive updating or overwriting characteristic of a tape recorder or a computer memory; for more on this argument, see Bjork, 1989).

The cue-dependent nature of retrieval assumed by the theory (not that that assumption is by any means unique to our theory) also has an important adaptive aspect. Normally inaccessible information becomes retrievable if the original environmental or situational cues to which that information was associated are reinstated. When we return to a town where we grew up, or to a school we attended, the environmental and interpersonal cues that are reinstated enable us to access memories that without those cues are not retrievable. That information tied to earlier contextual cues in our lives becomes more accessible when we are back in that context is, of course, adaptive: By virtue of our very return to an old context or situation, our need for access to information associated with that setting becomes greater.

Finally, although the argument is somewhat tenuous, the relative or absolute recovery of older memory representations with time may frequently be adaptive. Suppose that some new name or number or procedure has replaced an old name or number or procedure. With use of a new procedure, for example, that procedure will become more accessible in memory and the old procedure will become less accessible. The procedure in question, for example, might be the steps one goes through in moving blocks of text in a new word processing program. Suppose now that we stop using the new procedure. Many of the circumstances that would lead us to do so would also involve the old procedure becoming relevant again (we were on leave at some location where the new program was the institutional standard, and now we are returning to our own computer and office, or maybe we simply decided we did not like the new procedure). In such cases, the recovery of access to the old procedure, with disuse of the new procedure, would be adaptive.

ESTES' THEORY OF STIMULUS FLUCTUATION

Our theory of disuse represents our attempt to formulate, in a relatively neutral and abstract way, a representation of human memory that captures some storage and retrieval characteristics that we deem to be fundamental. As it turns out, a number of those "fundamental" characteristics are derivable from Estes' stimulus fluctuation theory. In this section we sketch the assumptions and achievements of the fluctuation theory.

In two papers published in the 1955 volume of the *Psychological Review*, William K. Estes shifted the focus of his "statistical theory of learning," from acquisition and extinction processes taking place within an experimental session, to between-session phenomena such as spontaneous recovery, forgetting, and distribution-of-practice effects on learning. The key idea was to assume that the stimulus aspects of a given environmental situation that were actually available

to (that is, could potentially have an influence on) the organism changed continuously over time. Such fluctuation processes could, then, result in a quite different functional stimulus situation when an organism returned to the same nominal experimental environment.

The initial assumptions of Estes' statistical learning theory, as put forward in his now classic 1950 paper, were summarized as follows by Estes (1955a) himself.

> a. Any environmental situation, as constituted at a given time, determines for a given organism a population of stimulus events from which a sample affects the organism's behavior at any instant; in statistical learning theories the population is conceptualized as a set of stimulus elements from which a random sample is drawn on each trial.
> b. Conditioning and extinction occur only with respect to the elements sampled on a trial.
> c. The behaviors available to an organism in a given situation may be categorized into mutually exclusive and exhaustive response classes.
> d. At any time, each stimulus element in the population is conditioned to exactly one of the response classes. (pp. 146–147)

Those initial assumptions presumed a perfectly constant physical environment in which the population of stimulus "elements" from which the organism sampled on any given trial was fixed. Such an idealized situation was unrealistic Estes argued, particularly if one wanted to understand recovery and forgetting across delays separating successive experimental sessions in the same or altered contexts. He assumed instead that at any moment only a subset of the possible stimulus elements in a given situation was available to the organism to sample, and that a random process over time governed whether a given stimulus element would be available or unavailable to the organism; that is, elements "fluctuated" between available and unavailable states. Such elements were presumed to correspond to "a large number of independently variable components or aspects of the environmental situation, all of which undergo constant random fluctuation" (1955a, p. 147). Bower and Hilgard (1981) summarized the possible reasons for such fluctuation as ". . . day to day fluctuation in the temperature or humidity of the experimental room, changes in the subject's internal milieu, postural set, or attitudes, sensitivity of various receptors, and so forth" (p. 227).

Achievements of the Stimulus Fluctuation Theory

The fluctuation idea implemented within the framework of stimulus sampling theory was able to account for a wide range of phenomena, in many cases quantitatively as well as qualitatively.

Spontaneous Recovery and Regression (Estes, 1955a). One of the elegant outcomes of Estes' analysis is that spontaneous recovery and forgetting, from the standpoint of the fluctuation theory, became formally equivalent. Assume that in an experimental apparatus of some kind an animal is reinforced for a certain behavior. At the beginning of acquisition, none of the elements, available or unavailable, would be conditioned to the appropriate response. At the end of a training session (given typical values of the fluctuation parameters), most or all of the elements in the available set will be conditioned, so the animal may be responding at a high rate, but most of the unavailable elements will remain unconditioned to the target response. If the animal is now given a period away from the apparatus, there will be forgetting; that is, a reduced rate of responding when the animal is brought back to the apparatus. Such forgetting is the consequence of conditioned elements fluctuating from the available to the unavailable set and unconditioned elements fluctuating from the unavailable to the available set. With a long enough delay, the system will reach an equilibrium state where the ratio of conditioned to unconditioned elements in the available and unavailable sets will be the same; that ratio will determine the asymptote of the forgetting function.

Assume now that as a consequence of repeated training sessions the animal is fully trained, which would mean that all the elements in the available *and* the unavailable sets are conditioned to the target response. An extinction session carried out with such a fully trained animal would correspond to unconditioning most or all of the elements in the available set. Many elements in the unavailable set would remain conditioned, however, and the fluctuation of those elements into the available set as a consequence of the animal being given a period away from the apparatus will lead to spontaneous recovery. The same type of dynamics, then, leads to both forgetting and spontaneous recovery.

The effects of overlearning and repeated learning on the rate of forgetting, and the effects of repeated extinction sessions on spontaneous recovery were accounted for in a natural way by the theory. The intervals between repeated learning sessions permit unconditioned elements in the unavailable set to fluctuate into the available set where they can be conditioned during the next training session. In an analogous fashion, the intervals between repeated extinction sessions permit conditioned elements to fluctuate from the unavailable to the available set, where they can then be unconditioned. The consequence of repeated learning or extinction sessions, then, is to condition (learning) or uncondition (extinction) more of the total population of stimulus elements, available and unavailable, thereby slowing the subsequent rate of forgetting or spontaneous recovery. The asymptotes of the forgetting and spontaneous-recovery functions will, of course, also be influenced by repeated learning or extinction sessions. In the limit, if all the elements in the available and unavailable sets are conditioned or unconditioned, respectively, there can be no forgetting or spontaneous recovery at all.

Distributional Phenomena in Learning (Estes, 1955b). In the second of his 1955 papers, Estes begins by saying "One aspect of distribution of practice that recent reviewers . . . seem agreed upon is its persistent refractoriness to any general theoretical interpretation" (p. 369). He then goes on, in our opinion, to provide such an interpretation—an interpretation that has been rediscovered in various guises periodically by subsequent researchers.

Within the framework of the fluctuation theory, Estes distinguished between *habit strength* and *response strength* (which correspond, as mentioned earlier, to storage strength and retrieval strength in our theory of disuse). Habit strength is determined by the total number of conditioned elements across the available and unavailable sets. Appropriate observable indices of habit strength are measures such as resistance to extinction and rate of forgetting. Response strength, on the other hand, is determined by the proportion of conditioned elements in the available set at any given point in time. It is reflected in terms of the momentary probability, rate, or latency of responding.

The importance of this distinction, with respect to the analysis of distribution (spacing) of practice, is that between-session stimulus fluctuation will have the opposite effect on the growth of habit strength and response strength. Increasing the spacing of trials or experimental sessions will increase stimulus fluctuation between trials or sessions, which will enhance the rate of growth of habit strength (because more new unconditioned elements will be available to condition during each trial or session) but will decrease the rate of growth of response strength, at least early in training, because more conditioned elements will fluctuate out of the available set from one trial or period to the next. Using the quantitative machinery of stimulus-sampling theory, Estes was able to account for a large variety of existing spacing effects in the conditioning and extinction literature. Among those phenomena are (a) that resistance to extinction is enhanced by spacing of conditioning periods; (b) that retention is increased by spacing learning periods with the advantage increasing with the length of the retention interval; and (c) that learning (or extinction) tends to be faster at the beginning of training with massed trials, but later in training the advantage shifts to spaced trials.

The observed interaction of spacing interval and retention interval (massing better for shorter retention intervals, spacing better for longer intervals) was later observed in experiments with humans as well (e.g., Glenberg, 1979; Glenberg & Lehmann, 1980; Peterson, Hillner, & Saltzman, 1962). That interaction, which emerges in a natural way from the fluctuation theory, has been particularly difficult for most other theories of the spacing effect to predict.

Finally, the fluctuation theory provides a staightforward interpretation for what seems to be a "paradoxical" aspect of the fact that long-term retention is enhanced by spaced repetitions. Melton (1967) pointed out that such an effect seems to suggest that forgetting helps memory, because more forgetting takes place prior to the second presentation the greater the interval from the first to the second presentation. Such "forgetting" in the fluctuation model corresponds to conditioned elements moving from the available to the unavailable set and unconditioned elements moving from the unavailable to the available set, which, of course, is exactly what is necessary to maximize the effectiveness of the second presentation.

Developments Since 1955. It is not possible here to track fully the influence of the fluctuation theory in the years following the two 1955 papers. A few examples, however, may illustrate the explanatory range of the fluctuation idea. Probably the most impressive extension of the theory has been to interference phenomena, particularly interference and recovery phenomena as a function of temporal variables and degree of learning. Estes (1959) presented some initial examples of how the fluctuation idea could account for certain retroactive and proactive effects, and Bower (1967) gave a more thorough and explicit treatment of interference phenomena in terms of fluctuation mechanisms. Recently, in an effort to account for a variety of interference and forgetting phenomena, Mensink and Raaijmakers (1988) incorporated the contextual fluctuation mechanism of stimulus sampling theory within the SAM framework of Raaijmakers and Shiffrin (1981). As we read that impressive paper, the successful analysis of various classic interference phenomena are more a consequence of the contextual-fluctuation assumption than of the SAM framework itself.

The fluctuation idea also has been extended to account for tests as learning and forgetting-prevention events (Izawa, 1971; Whitten & Bjork, 1977), and to experiments on verbal short-term memory (Bower, 1972; Estes, 1971; Peterson, 1963). In a particularly systematic and thorough exercise, Bower (1972) showed how encoding variability ideas, implemented within the stimulus-fluctuation version of stimulus sampling theory, can account for characteristics of recognition, temporal-lag judgments, list-differentiation, and consolidation/retrograde-amnesia.

Beyond such explicit, quantitative, extrapolations of the fluctuation theory, there have been widespread conceptual influences of the fluctuation idea. The fluctuation-plus-sampling process is a forerunner of the various incarnations of the encoding-variability idea (e.g., Martin, 1968). Glenberg's (1979) ambitious theory of spacing phenomena in human memory, in particular, can be viewed as a generalization and extension of Estes' original stimulus-fluctuation idea. In some cases, researchers have apparently not been aware that their ideas were expressed earlier in stimulus-fluctuation terms. Cuddy and Jacoby's notion (1982; Jacoby, 1978), for example, that for a repetition to be effective there needs to be some forgetting of the initial presentation, is stated explicitly by Estes (as mentioned earlier) in the second of the 1955 papers. (In fairness to Cuddy and Jacoby, their argument as to *why* such forgetting is helpful differs from the fluctuation-theory argument.)

Fluctuation Theory as a Guide to Intuition: An Illustration

One informal measure of a theory's value is whether it can serve as a guide to one's intuition, in the sense that it is a better basis for prediction than one's "common sense." Put in a way that seems to be gaining popularity, does the theory pass the "grandmother test"; that is, does it tell us something more than our grandmother could have told us? On that measure, fluctuation theory holds up very well indeed.

To illustrate that unintuitive implications of the fluctuation theory often prove to be correct, we describe briefly the results of an unpublished experiment carried out by the first author some years ago.[1] In this study, 60 subjects each went through a deck of 88 cards at a 10-s rate. On a given card in the deck, a pair of 4-letter nouns was presented for study (e.g., FERN-WASP), or the left hand member of a prior pair was presented as a test of the subjects' memory for the right hand member of that pair (e.g., ROPE- ??). Thus, the experiment employed a continuous paired-associate paradigm, but there was an updating requirement: About half the time, the first study card involving a given stimulus (FERN-WASP) was followed, after a variable number of intervening cards involving other pairs, by a second study card on which that stimulus was paired with a new response word (e.g., FERN-PULP). On the later test of that pair in the deck (FERN- ??), subjects were to give only the current, most recent, response (PULP). Subjects could not ignore the first study card involving a given stimulus, however, because half the time, in unpredictable fashion, that first pairing was the only study card involving that stimulus (that is, there was no second response paired with that stimulus) in which case the correct response to the test (FERN-??) was WASP.

The pairs involving a repeated stimulus, then, corresponded to an AB-AC paradigm. The question motivating the experiment was a simple one: What would happen to the recall of the second response (C) as a function of the interval between the AB and AC pairings? And would the effects of the AB-AC interval interact with the retention interval from AC to the test? In the experimental design those two intervals were co-varied: 0, 3, or 9 cards intervened between the AB and AC cards, and 3 or 9 cards intervened between the AC pairing and the test. Across those six conditions, the interval from the AB pairing to the test, then, was 3, 6, 9, 12, and 18 events involving other pairs. For the purpose of providing a baseline against which to evaluate the intrusion rates of the B response in the different AB-AC conditions, a subset of the once-presented AB pairs was tested at those same AB-Test delays (where, now, B is the correct response).

[1]Lisa White and Brenda Hasson provided invaluable assistance in this research.

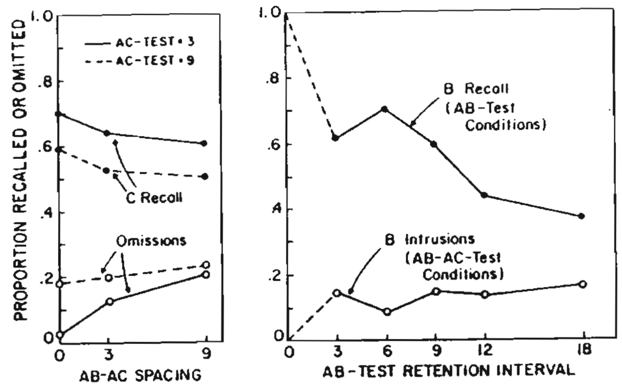

FIG. 2.2. Proportion of responses that were correct (C recall) or omissions as a function of the AB-AC spacing interval and the AC-Test retention interval (left panel); proportions of B responses correctly recalled and intruded, respectively, as a function of the AB-Test retention interval in the control (no AC presentation) and AB-AC conditions.

From a common-sense standpoint, there seems little doubt (for a fixed retention interval) that the recall of C should increase with the AB-AC interval. After all, not only should the AC event be more discriminable from the AB event the greater the spacing interval, but the B response should be weaker as well (because the interval from AB to the test increases with the AB-AC interval). About the only reason to expect otherwise—and the reason we bothered to do the experiment—is the spacing effect itself. If one views the repetition case (AB-AB-Test) as a positive-transfer paradigm, then it could be that the conditions that maximize positive transfer (increased AB-AB spacing) might also maximize negative transfer in the AB-AC paradigm.

The results of the experiment are shown in Fig. 2.2. In the left panel, proportion correct recall of C responses and omissions are plotted as a function of AB-AC spacing for each retention interval. In the right panel, recall of the B response from the AB once-presented control items, and intrusions of the B response in the AB-AC conditions, are plotted as a function of the AB-Test retention interval.

Common sense notwithstanding, the level of C recall in the left panel goes down (and omissions go up) as a function of the AB-AC spacing interval. Performance is also higher at the shorter AC-Test retention interval, which is no surprise. In the right panel, recall in the AB control condition goes down with retention interval, as it should, but the rate of intrusions of B in the AB-AC conditions does *not* show a corresponding decrease across the same AB-Test intervals.

The fluctuation theory has little difficulty in accounting for the general pattern of results in Fig. 2.2. The basic argument is that the longer the interval from AB

to AC the more opportunity there is for elements "conditioned" to B to fluctuate to the unavailable state prior to the AC event. To the degree they are lurking in the unavailable set at the time the AC pairing is studied, they are protected from being "counterconditioned" to the C response, which means they may fluctuate back to the available set during the AC-Test interval. If they are in the available set at the time of test, they may cause intrusions of the B response (right panel) or they may contribute to omissions (left panel) by, for example, leading to a sample of elements that does not favor the B or the C response.

Some Metacomments on the 1955 Papers

We cannot resist sharing a few reactions we had in rereading the original 1955 papers. First of all, those papers are surprisingly short; they are nine and one-half and nine pages long, respectively, and those are the small pages of 1955, not the current large pages.

We were also surprised in going back over those papers how much the empirical phenomena Estes cited, and the terminology he used, were drawn from animal research. In virtually every case where results from the animal literature are cited, analogous human results could have been cited instead or in addition. Those relevant human results, in fact, were nicely summarized in the outstanding books on experimental psychology published by Osgood (1953) and Woodworth and Schlosberg (1954) just prior to the Estes 1955 papers.

The only basis for our surprise, of course, is that *we* think in terms of the more cognitive Bill Estes who emerged in the late 1950s, and blossomed in the early 1960s. According to Estes' (1989) autobiography, the influence of his graduate mentor, B. F. Skinner—who would not let him use the word "memory" in their coauthored papers—and an early philosophical commitment to operationism guided his empirical and theoretical work during the first 15 to 20 years of his professional career. His eventual shift toward a rigorous type of constructivism was a consequence, in part, of the influence of Patrick Suppes, Gordon Bower, Richard Atkinson, and other neomentalists at Stanford University. (The use of "neomentalists" is our, not Estes', characterization.)

From a philosophy-of-science standpoint, the 1955 papers seem to constitute a kind of balancing act. There seems to be an effort to defend such theorizing to the radical behaviorists while rejecting a certain type of untestable mentalism.

All such considerations aside, we were struck again at how beautifully written and elegantly argued those papers are. In both papers, the ideas are sketched in very understandable terms before being phrased, in separate sections, in mathematical terms (". . . to permit self-selection diets for readers of varying mathematical appetites" [p. 369]). Even abstract points are stated in ways that are lively (". . . but an unfilled interval never remains permanently satisfying as an explanatory variable" [p. 145]) and entertaining ("Few hypothetical entities are so ill-favored that once having secured a foothold they cannot face out each new turn of empirical events with the aid of a few ad hoc assumptions" [p. 145]).

A COMPARISON OF FLUCTUATION THEORY AND OUR NEW THEORY OF DISUSE

On several occasions in this chapter, we have referred to our new theory of disuse as a "framework." Because there is a general correspondence between Estes's response strength and habit strength concepts to our retrieval strength and storage strength concepts, respectively, and because the theories account for many of the same phenomena, the question arises as to what extent the fluctuation theory has the storage and retrieval properties that define our theory of disuse. That is, to what degree does the fluctuation theory fit our framework?

Before we address that question, it is important to clarify that it is the *formal* properties of the two theories that are at issue. The original stimulus-response phrasing of the stimulus fluctuation theory, and its focus on conditioning and extinction phenomena in animals, makes the theory seem different in kind than our more abstract and human-memory-oriented theory of disuse. The basic fluctuation idea, however, can be expressed in more modern-sounding and cognitive ways. In Bower's (1972) representation, for example, contextual influences of various kinds are presumed to influence what subset of the set of possible encoding operations are activated when a given nominal stimulus is presented, and each resulting stimulus encoding is presumed to be associated with one or more responses in memory. The strength of a given associated response in memory, then, is equal to the proportion of active encodings to which it is associated.

Bower's (1972) assumption that more than one response can be associated to a given stimulus encoding violates an original assumption of stimulus-sampling theory (that one and only one response is conditioned to a given stimulus element). Such a change was necessary, according to Bower, given the ". . . evidence showing that learning a second association to a stimulus need not cause unlearning of an earlier association but rather only edited differentiation or temporary suppression of the earlier response" (p. 91).

Going further in that direction, one could assume that as a consequence of environmental/social/physiological/cognitive/ and emotional aspects of context, a single functional encoding of the nominal cue configuration is extracted by the human subject, which then defines a set of items in memory associated to that functional cue. The "response strength" of a given item to that cue would then be defined in terms of the recency or frequency of the past co-occurrences of that cue and response (*apple* would have greater response strength than *pomegranate* to the cue *fruit* because of the prior history of the organism). Habit strength would then correspond to the strength of association summed across all possible

functional encodings of a given nominal cue configuration. Available and unavailable response items in memory would then correspond to the set associated to the currently active encoding of that cue configuration and the sets associated to the other possible encodings of that configuration, respectively. What would fluctuate across successive presentations of a given nominal cue is the particular encoding or interpretation the subject extracts from or gives to that nominal cue. Such fluctuation would be the consequence of changes in one or more of the dimensions of context listed above.

The foregoing rephrasing of fluctuation theory is intended to illustrate that the theory can be phrased in ways that increase its apparent compatibility with the human-memory domain that is the focus of our theory of disuse. Most, but not all, of the formal properties of the fluctuation theory are left unchanged by such a reformulation. Before we comment on what properties might change, we first want to compare the original version of fluctuation theory to our theory of disuse.

Response Strength Versus Retrieval Strength

Similarities. Response strength, defined as the proportion of elements in the available set conditioned to a given response, corresponds in spirit and function to our retrieval-strength concept. Some formal similarities are the following:

1. The probability that a given response is accessible in a given situation is solely a function of its current response strength.
2. The increment in the strength of a given response as a consequence of a study (reinforcement) trial is a negatively accelerated function of that response's prior strength.
3. Following a study (training) session during which a given response is reinforced, the loss of response strength over time (forgetting) proceeds more slowly the higher the level of habit strength at the end of training.
4. The decrement in the response strength of a given response as a consequence of the reinforcement of a competing response is greater the higher the response strength of the response in question.
5. In general, although the immediate size of such decrements in response strength will depend only on the response strength of the response in question, not on its habit strength, the observable decrement in response strength over time will be smaller the greater its habit strength. (In the fluctuation model, of course, there need not be any actual decrement at all, as in the case of spontaneous recovery.)

Differences. The response-strength properties of the original fluctuation theory differ from the corresponding retrieval-strength properties in the following ways.

1. Without there being a reinforcement or feedback event, response strengths are unchanged as a consequence of the organism making a given response. This difference between the fluctuation and disuse theories is basic, because the act of retrieval is assumed to be a potent event (in altering the strength of the retrieved response—and the strengths of competing responses) in the theory of disuse.

2. Another significant difference is that response strength declines over time not to an eventual asymptote of zero, as in the theory of disuse, but to an asymptote defined by habit strength. That is, response strength declines to an asymptote equal to the overall proportion of conditioned elements in the available and unavailable sets (habit strength).

3. The global limit on retrieval capacity assumed in the theory of disuse has no counterpart in the fluctuation theory. There is no effect on a given response's strength in one stimulus situation as a consequence of increasing a different response's strength (or even the same response's strength) in a different (non-overlapping) stimulus situation.

4. Finally, there are differences in the two theories as to how response/retrieval strength translates into response probability. In the fluctuation theory the translation is quite straightforward: Response strength is equal to the proportion of conditioned elements in the available set and response probability is equal to the proportion of conditioned elements in a given sample from the available set. The expected value of the latter proportion is, of course, equal to the former proportion; in virtually all applications of the model the variance in momentary response probability owing to the sampling process can be ignored. In the theory of disuse, at the level it is specified in this chapter, response probability depends on both relative and absolute retrieval strengths: A given item must be discriminated from the other items associated to a given cue configuration, which is a function of its strength relative to the strengths of the other cued items, and it must be reconstructed once it is discriminated, which is a function of its absolute retrieval strength. This two-step process, analogous to the sampling-plus-recovery process in Raaijmakers and Shiffrin's (1981) SAM model, attributes a level of complexity to the retrieval process that does not seem compatible with the stimulus-response roots of stimulus-sampling theory.

Habit Strength Versus Storage Strength

Similarities. Habit strength, defined as the overall proportion of elements in the available and unavailable sets conditioned to a given response, corresponds in a general way to storage strength in the theory of disuse. Some particular similarities are the following:

1. Habit strength, like storage strength, has no direct effects on performance. Its influences on performance are manifested through its effects on the gain or loss of response strengths across study sessions or intervening intervals, respectively. Those effects on response strength, as cited earlier, are quite similar to the presumed influences of storage strength on the gain or loss of retrieval strength.

2. In the fluctuation model, as in the disuse model, there is no limit on learning (storage) as expressed by habit strength summed across all possible cue/stimulus configurations and associated responses.

3. The increment in habit strength as a consequence of a reinforcement or study trial is a negatively accelerated function of current response strength.

4. In general, the increment in habit strength is also a decreasing function of current habit strength. (In a given situation, however, the effect of a study trial on habit strength will depend on how elements already conditioned to the target response are partitioned between the available and unavailable sets: The fewer conditioned elements there are in the available set, the larger the increment in habit strength.)

Differences. The following are the most significant differences in the properties assigned to habit strength and storage strength.

1. A major difference in the two theories is that storage strength, once accumulated, is never lost, whereas habit strength is reduced in the fluctuation model whenever elements conditioned to that target response are reconditioned to another response. Habit strength can go to zero in the fluctuation model, as would be the case after repeated spaced extinction sessions. (From such a point of "complete extinction," relearning the extinguished response should be as slow as the original learning of that response, whereas in the theory of disuse, relearning would be much more rapid than original learning, owing to the storage strength remaining from the original learning phase.) Another way to phrase this difference between the theories is to say that changes in response strength owing to conditioning or unconditioning of stimulus elements in the available set result in corresponding changes in habit strength, whereas there is no such correspondence in the theory of disuse. Changes in response strength that result from the fluctuation process, on the other hand, do leave habit strength unchanged.

2. Although there is no limit on habit strengths summed across situations and responses, there is a limit on the habit strengths summed across the responses associated to a given stimulus situation. Increasing the habit strength of one response decreases, in general, the habit strength of other responses (there are only so many stimulus elements corresponding to a given configuration of cues, and only one response can be conditioned to a given element). In the disuse theory, on the other hand, increasing the storage strength of one member in the set of items associated to a given cue configuration does not alter the storage strengths of the other items in that set.

Towards a Resolution of Certain Differences

Some of the differences between the formal properties of the original version of the fluctuation theory and the properties of our new theory of disuse do not seem fundamental to the fluctuation idea, per se, but to the initial phrasing of certain assumptions. We indicate in the following sections how changes in those assumptions might reduce or eliminate some of the differences outlined earlier.

Test Effects. As originally formulated, the production of a response, without some kind of feedback, does not alter response strengths or habit strengths in the fluctuation model. The potent effects of a test event can, however, be accommodated in several ways in the fluctuation theory. Izawa (1971) was able to account quantitatively for the effects of unreinforced test cycles in paired-associate learning with a modified version of the fluctuation model. She was able to capture two distinct benefits of test trials (they retard forgetting and they potentiate subsequent study trials) by assuming that a test cycle perturbs the fluctuation process (conditioned elements in the available set remain available, unconditioned elements in the available set fluctuate to the unavailable set with probability k, and elements in the unavailable set fluctuate to the available set with probability k'). Whitten and Bjork (1977) were able to fit the quantitative consequences of tests in a modified Brown–Peterson distractor paradigm by assuming that a successful recall slowed the fluctuation rate of alternative "retrieval routes" to a given item in memory. Such an assumption has the effect of increasing the likelihood that a successful retrieval route will again be available in the future.

In the Izawa (1971) and Whitten and Bjork (1977) formulations, no learning is presumed to take place on test trials per se. One could assume that sampled elements conditioned to responses other than the one produced are reconditioned to the response given. Such an assumption, however, yields an increased rate of responding across a series of unreinforced test trials. Such "hypermnesia" (Erdelyi & Becker, 1974) does occur in well defined situations (see, e.g., Roediger & Challis, 1989), but such cases are the exception, not the rule.

The Asymptote of Forgetting. In the fluctuation model, response strength does not decrease to an eventual asymptote of zero, as is assumed in our new theory of disuse, but to a level that corresponds to habit strength. Response strength cannot be driven to zero unless habit strength, via repeated extinction sessions, is driven to zero. To the extent, however, that (a) little or no fluctuation takes place during a training session, and (b) the size of the unavailable set is large with respect to the size of the available set, the asymptote of forgetting will be close to zero. (Researchers have typically made the former assumption, and best-fitting parameter estimates have typically implied the latter relationship.) Under those conditions, performance can reach a high level at the end of training, but the overall proportion of conditioned elements across the available and unavailable sets will remain small.

The Loss of Habit Strength. In the fluctuation model, any loss of response strength (by virtue of elements in the available set being reconditioned to another response) results in a corresponding decrease in habit strength. In the theory of disuse, the storage strength of a given response is not altered by changes in the retrieval strength (or storage strength) of other items. Put another way, a response's habit strength in the fluctuation model is tied to its response strength, and to the response and habit strengths of other responses, in ways that storage strength is not in the theory of disuse. There is only so much habit strength and response strength to go around in the fluctuation model because the number of elements is assumed to be fixed and each element is assumed to be associated to one and only one response. One consequence of this property of the fluctuation theory, among others, is that the habit strength of a given response can be driven to zero via extensive training on another response (or responses), however high that response habit strength may have been at one point.

In Bower's (1972) formulation, in which more than one response can be associated to a given stimulus encoding, items can gain in response or habit strength without altering the response strength or habit strength of other items. In Bower's version of the theory, an item's response strength (defined as the proportion of active encodings to which it is associated), or its habit strength (defined as the proportion of all possible encodings to which it is associated), can be altered without necessarily changing the response or habit strengths of other items.

There is, however, a cost associated with the flexibility that is gained in Bower's (1972) reformulation of the theory: The translation of response strength into response probabilities becomes complicated. Using the simple proportionality rule in the original formulation of the fluctuation model (that is, setting the probability of a given response equal to the proportion of elements in the available set conditioned to that response), the sum of the probabilities of alternative responses can, in Bower's formulation, exceed one (in fact, in principle, the probability of every response could be unity). Assuming it is the case that a variety of responses (items in memory) are triggered (associated to) a given cue, rules must then be specified to designate which response is given (or in what order multiple responses are given)—which is what Bower does in his paper for each of a set of experimental paradigms.

Remaining Differences

In one sense, it is a mistake to overemphasize the formal properties that correspond, or could be made to correspond, in the fluctuation and disuse theories. There *are* some fundamental conceptual differences between the two theories. In particular, what is meant by retrieval strength in our theory and what is meant by response strength in Estes's fluctuation theory differ in a fundamental way. We think of retrieval strength as a property of an item's representation in memory:

how active or primed that representation is at a given point in time. That is, retrieval strength is a property of the item per se. Response strength, on the other hand, is defined in terms of the association of a given response to aspects of a stimulus situation that happens to be sampled by the organism. As long as the organism has no intervening experience in that situation, response strengths stay intact however long the intervening interval; whereas retrieval strength, as an activity attribute of an item in memory, decreases as a consequence of intervening activities whatever the setting of those activities.

One other difference seems significant from a conceptual, if not formal, standpoint. In the fluctuation theory, the sampling dynamics are formulated with respect to the stimulus constellation impinging on the organism. Retrieving a response once stimulus sampling has taken place is a relatively automatic (though probabilistic) process governed by stimulus-response association. In our theory of disuse, on the other hand, the "sampling" dynamics are on the contents of the memory system itself. The configuration of cues presented to the subject is assumed to restrict (relatively automatically) the search of memory to an associated set of items in memory. Retrieval of a particular item in that set requires that it be discriminated from the other items in the set (a function of its relative retrieval strength) and then reconstructed (a function of its absolute retrieval strength).

CONCLUDING COMMENTS

The new theory of disuse represents our conjecture as to the adaptive interplay of storage and retrieval processes in human memory. In formulating the theory, we have taken as an article of faith that what appear to be peculiarities of human memory, as exhibited in certain real-world and laboratory settings, are in fact reflections of storage and retrieval processes that are adaptive in the overall functioning of human memory.

The theory as formulated, in the distinction it makes between storage strength and retrieval strength, and in many of its formal properties, bears a close relationship to the stimulus-fluctuation theory published by William K. Estes (1955a, 1955b) 35 years ago. As we have tried to indicate, the fluctuation idea remains one of the major theoretical insights in the history of research on learning and memory. It is one—but only one—of the significant contributions to our field by the man in whose honor we write this chapter.

ACKNOWLEDGMENTS

The preparation of this report was supported in part by Grants 4-564040-RB-19900 and 4-564040-EB-19900 to the first and second authors, respectively, from the Committee on Research, University of California. We thank Steve Clark, Harold Gelfand, Todd Gross, Barbara Spellman, and, especially, Michael Anderson, for their criticisms of and contributions to our theory of disuse.

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. (1989). A rational analysis of human memory. In H. L. Roediger III & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 195–210). Hillsdale, NJ: Lawrence Erlbaum Associates.

Anderson, M. C., & Bjork, R. A. (1990, November). *Category-based retrieval inhibition in human memory*. Paper presented at the meeting of the Psychonomic Society, New Orleans, LA.

Baddeley, A. (1988). But what the hell's it for? In M. M. Gruneberg, P. E. Morris, & R. N. Sykes (Eds.), *Practical aspects of memory: Current research and issues: Vol. 1. Memory in everyday life* (pp. 1–15). London: Wiley.

Bartlett, J. C. (1977). Effects of immediate testing on delayed retrieval: Search and recovery operations with four types of cue. *Journal of Experimental Psychology: Human Learning and Memory, 3*, 719–732.

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

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

Bjork, R. A. (1989). Retrieval inhibition as an adaptive mechanism in human memory. In H. L. Roediger III & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 195–210). Hillsdale, NJ: Lawrence Erlbaum Associates.

Bjork, E. L., & Bjork, R. A. (1988). On the adaptive aspects of retrieval failure in autobiographical memory. In M. M. Gruneberg, P. E. Morris, & R. N. Sykes (Eds.), *Practical aspects of memory: Current research and issues: Vol. 1. Memory in everyday life* (pp. 283–288). London: Wiley.

Bjork, R. A., & Geiselman, R. E. (1978). Constituent processes in the differentiation of items in memory. *Journal of Experimental Psychology: Human Learning and Memory, 4*, 344–361.

Bjork, R. A., & Richardson-Klavehn, A. (1989). On the puzzling relationship between environmental context and human memory. In C. Izawa (Ed.), *Current issues in cognitive processes: The Tulane Flowerree symposium on cognition* (pp. 313–344). Hillsdale, NJ: Lawrence Erlbaum Associates.

Bjork, R. A., & Whitten, W. B. (1974). Recency-sensitive retrieval processes in long-term free recall. *Cognitive Psychology, 6*, 173–189.

Bower, G. H. (1967). Verbal learning. In H. H. Helson & W. Bevan (Eds.), *Contemporary approaches to psychology*. Princeton, NJ: Van Nostrand.

Bower, G. H. (1972). Stimulus-sampling theory of encoding variability. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory*. Washington, DC: W. H. Winston & Sons.

Bower, G. H., & Hilgard, E. R. (1981). *Theories of learning* (5th ed.). Englewood Cliffs, NJ: Prentice-Hall.

Bower, G. H., & Reitman, J. S. (1972). Mnemonic elaboration in multilist learning. *Journal of Verbal Learning and Verbal Behavior, 11*, 478–485.

Briggs, G. E. (1954). Acquisition, extinction and recovery functions in retroactive inhibition. *Journal of Experimental Psychology, 47*, 285–293.

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.
Cuddy, L. J., & Jacoby, L. L. (1982). When forgetting helps. *Verbal Learning and Verbal Behavior, 21,* 451–467.
Ebbinghaus, H. (1964). *Memory: A contribution to experimental psychology.* (H. A. Ruger & C. E. Bussenius, translators) New York: Dover. (Original work published 1885).
Erdelyi, M. H., & Becker, J. (1974). Hypermnesia for pictures: Incremental memory for pictures but not for words in multiple recall trials. *Cognitive Psychology, 6,* 159–171.
Estes, W. K. (1950). Toward a statistical theory of learning. *Psychological Review, 57,* 94–107.
Estes, W. K. (1955a). Statistical theory of spontaneous recovery and regression. *Psychological Review, 62,* 145–154.
Estes, W. K. (1955b). Statistical theory of distributional phenomena in learning. *Psychological Review, 62,* 369–377.
Estes, W. K. (1959). The statistical approach to learning theory. In S. Koch (Ed.), *Psychology: A study of a science* (Vol. 2, pp. 380–491). New York: McGraw-Hill.
Estes, W. K. (1971). Learning and memory. In E. F. Beckenbach & C. B. Tomkins (Eds.), *Concepts of communication* (pp. 282–300). New York: Wiley.
Estes, W. K. (1989). William K. Estes. In G. Lindzey (Ed.), *A history of psychology in autobiography* (pp. 95–124). Stanford, CA: Stanford University Press.
Forrester, W. E. (1970). Retroactive inhibition and spontaneous recovery in the A-B, D-C paradigm. *Journal of Verbal Learning and Verbal Behavior, 9,* 525–528.
Gardiner, J. M., Craik, F. I. M., & Bleasdale, F. A. (1973). Retrieval difficulty and subsequent recall. *Memory & Cognition, 1,* 213–216.
Glenberg, A. M. (1979). Component-levels theory of the effects of spacing of repetitions on recall and recognition. *Memory & Cognition, 7,* 95–112.
Glenberg, A. M., & Lehmann, T. S. (1980). Spacing repetitions over 1 week. *Memory & Cognition, 8,* 528–538.
Hogan, R. M., & Kintsch, W. (1971). Differential effects of study and test trials on long-term recognition and recall. *Journal of Verbal Learning and Verbal Behavior, 10,* 562–567.
Hull, C. L. (1943). *The principles of behavior.* New York: Appleton-Century-Crofts.
Izawa, C. (1970). Optimal potentiating effects and forgetting-prevention effects of tests in paired-associate learning. *Journal of Experimental Psychology, 83,* 340–344.
Izawa, C. (1971). The test trial potentiating model. *Journal of Mathematical Psychology, 8,* 220–224.
Jacoby, L. L. (1978). On interpreting the effects of repetition: Solving a problem versus remembering a solution. *Journal of Verbal Learning and Verbal Behavior, 17,* 649–667.
Jost, A. (1897). Die Assoziationfestigkeit in Ihrer Abhangigheit von der Verteiling der Wiederholungen. *Zeitschrift fur Psychologie, 14,* 436–472.
Koppenaal, R. J. (1963). Time changes in strength of A-B, A-C lists: Spontaneous recovery? *Journal of Verbal Learning and Verbal Behavior, 2,* 310–319.
Krueger, W. C. F. (1929). The effects of overlearning on retention. *Journal of Experimental Psychology, 12,* 71–78.
Landauer, T. K. (1986). How much do people remember? *Cognitive Science, 10,* 477–493.
Landauer, T. K., & Bjork, R. A. (1978). Optimal rehearsal patterns and name learning. In M. M. Gruneberg, P. E. Morris, & R. N. Sykes (Eds.), *Practical aspects of memory* (pp. 625–632). London: Academic Press.
Martin, E. (1968). Stimulus meaningfulness and paired-associate transfer: An encoding variability hypothesis. *Psychological Review, 75,* 421–441.
McGeoch, J. A. (1932). Forgetting and the law of disuse. *Psychological Review, 39,* 352–370.
Melton, A. W. (1967). Repetition and retrieval from memory. *Science, 158,* 532.
Melton, A. W., & Irwin, J. M. (1940). The influence of degree of interpolated learning on retroactive inhibition and the overt transfer of specific responses. *American Journal of Psychology, 53,* 173–203.

Mensink, G. J., & Raaijmakers, J. G. W. (1988). A model for interference and forgetting. *Psychological Review, 93,* 434–455.
Miller, G. A., & Gildea, R. M. (1987, September). How children learn words. *Scientific American,* pp. 94–99.
Modigliani, V. (1976). Effects on a later recall by delaying initial recall. *Journal of Experimental Psychology: Human Learning and Memory, 2,* 609–622.
Osgood, C. E. (1953). *Method and theory in experimental psychology.* New York: Oxford University Press.
Peterson, L. R. (1963). Immediate memory: Data and theory. In C. N. Cofer & B. S. Musgrave (Eds.), *Verbal behavior and verbal learning* (pp. 336–353). New York: McGraw-Hill.
Peterson, L. R., Hillner, K., & Saltzman, D. (1962). Time between pairings and short-term retention. *Journal of Experimental Psychology, 64,* 550–551.
Postman, L., Stark, K., & Fraser, J. (1968). Temporal changes in interference. *Journal of Verbal Learning and Verbal Behavior, 7,* 672–694.
Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review, 88,* 93–134.
Ratcliff, R. C., Clark, S., & Shiffrin, R. M. (1990). The list-strength effect: I. Data and discussion. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 163–178.
Richardson-Klavehn, A. (1988). *Effects of incidental environmental context on human memory: Elusive or nonexistent?* Unpublished doctoral dissertation, University of California, Los Angeles.
Roediger, H. L., & Challis, B. H. (1989). Hypermnesia: Improvements in recall with repeated testing. In H. L. Roediger III & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 175–199). Hillsdale, NJ: Lawrence Erlbaum Associates.
Santiago, H. C., & Wright, A. A. (1984). Pigeon memory: Same/different concept learning, serial probe recognition acquisition and probe delay effects in the serial position function. *Journal of Experimental Psychology: Animal Behavior Processes, 10,* 498–512.
Skinner, B. F. (1938). *The behavior of organisms.* New York: Appleton-Century-Crofts.
Thompson, C. P., Wenger, S. K., & Bartling, C. A. (1978). How recall facilitates subsequent recall: A reappraisal. *Journal of Experimental Psychology: Human Learning and Memory, 4,* 210–221.
Thorndike, E. L. (1914). *The psychology of learning.* New York: Teachers College.
Tulving, E., & Thompson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review, 80,* 352–373.
Whitten, W. B. II, & Bjork, R. A. (1977). Learning from tests: Effects of spacing. *Journal of Verbal Learning and Verbal Behavior, 16,* 465–478.
Woodworth, R. S., & Schlosberg, H. (1954). *Experimental psychology.* New York: Holt.
Wright, A. A. (1989). Memory processing by pigeons, monkeys, and people. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 24, pp. 25–70). New York: Academic Press.
Wright, A. A., Santiago, H. C., & Sands, S. F. (1984). Monkey memory: Same/different concept learning, serial probe acquisition, and probe delay effects. *Journal of Experimental Psychology: Animal Behavior Processes, 10,* 513–529.
Wright, A. A., Santiago, H. C., Sands, S. F., Kendrick, D. F., & Cook, R. G. (1985). Memory processing of serial lists by pigeons, monkeys, and people. *Science, 229,* 287–289.