1 Forgetting

Preliminary considerations

Henry L. Roediger III,
Yana Weinstein, and Pooja K. Agarwal
Washington University in St. Louis, USA

> The existence of forgetting has never been proved: we only know that some things do not come to our mind when we want them to.
>
> (Friedrich Nietzsche, 1844–1900)

Of all the common afflictions from which humankind suffers, forgetting is probably the most common. Each of us, every day, forgets something we wish we could remember. It might be something we have done, something we intended to do, a fact, a name of a person or restaurant, and so on ad infinitum. As we age, our incidents of forgetting increase and we worry more about them. A whole industry of books, tapes, and even new mental gymnasia has grown up to deal with the cognitive frailties of old age, the primary one being rampant forgetting. Compared to other nuisances of life, forgetting probably tops the list. The "common cold" is actually quite rare compared to forgetting in all its manifestations. As Underwood (1966) wrote: "Forgetting is a most exasperating and sometimes even painful phenomenon" (p. 542). More recently, Nairne and Pandeirada (2008) maintained that for most people "forgetting is a scourge, a nuisance, a breakdown in an otherwise efficient mental capacity" (p. 179), although they quickly noted that there is often an adaptive value in forgetting too.

Despite the fact that psychologists have been studying learning and memory for 125 years, the current volume is the only one we can find devoted solely to the topic of forgetting. "Forgetting" is a term used in the titles of many works of fiction and even cultural critique (see Markowitsch & Brand, Chapter 2), but this volume is the first scientific one devoted to it. Strange, you might think.

Given the ubiquity of forgetting in our daily lives, the quote by Nietzsche that heads our chapter must seem stranger still. Given its ubiquity, how can the existence of forgetting be doubted? Difficulties of these sorts usually revolve around matters of definition, and that is the case here. We turn to this issue first.

Defining forgetting

According to the authors of the *International encyclopedia of the social sciences*: "It seems quite unnecessary to be concerned with a definition of 'forgetting' " (Sills & Merton, 1968, p. 536). Nonetheless, psychologists have attempted to define forgetting in several different ways. Cubelli (Chapter 3) provides a thorough exploration of the various extant definitions of forgetting, and below we give a general overview. Before undertaking the task of examining these issues, however, we review some preliminary considerations. At least since Köhler (1947, p. 279), psychologists have found it useful to distinguish among three stages in the learning/memory process: acquisition (encoding), storage (maintenance or persistence), and retrieval (utilization of stored information, see too Melton, 1963; Weiner, 1966). Encoding or acquisition is the initial process in learning, although this process may be extended in time as a memory trace (a persisting representation) formed through consolidation. Only events that have been securely encoded or learned in the first place can be said to be forgotten; it makes no sense to say that one has forgotten the 15th name in the Auckland, NZ, telephone book or the capital of Mars, because one never knew these bits of information in the first place. We take Tulving's definition of forgetting – "the inability to recall something now that could be recalled on an earlier occasion" (1974, p. 74) – as our starting point in considering more complex definitions. We consider first the strongest form of the concept of forgetting, the one implicit in the quote from Nietzsche.

Forgetting as complete loss from storage

Davis (2008) defines the strong form of forgetting as "the theoretical possibility that refers to a total erasure of the original memory that cannot be recalled, no matter what techniques are used to aid recall" (p. 317). Given the context of his chapter, we feel sure he would be willing to include not just measures of recall, but any measure (explicit or implicit, direct or indirect) of the prior experience having been encoded in the nervous system. Davis argued that it would only be possible to look for "strong" forgetting in simple organisms (e.g., simple gastropods like slugs) where the entire neural circuitry has been mapped out. "Only when all the cellular and molecular events that occur when a memory is formed return to their original state would I say this would be evidence for true forgetting" (Davis, 2008, p. 317).

To our knowledge, no evidence for this strong form of forgetting has been produced even in simpler organisms; and since all the research in the present volume is about forgetting in organisms more complex than mollusks, it would be practically impossible to obtain evidence for this strong form of forgetting. Even if every test known to psychologists failed to show evidence for any sort of trace of past experience, the possibility remains that a change owing to that prior experience (some latent memory trace) still remains.

Davis (2008) concluded that the strong form of forgetting is not scientifically useful, and we agree with him. We can ask the further question: If the strong form of forgetting can never be proved (as Nietzche's dictum states), does this mean that forgetting in this sense never occurs? We think the answer to this question must be no (although we cannot prove it). Think of all the events and happenings that occurred to you when you were 7 years old, ones you could have easily reported the next day (so they were encoded). Do you still really have traces of all these events lying dormant in your brain, waiting for the right cue to become active again? We strongly doubt it. Probably the many of the millions of events, conversations, facts, people, and so on that are encountered in everyday life and at one point committed to memory do suffer the strong form of forgetting by being obliterated from our nervous systems. However, that is a matter of faith, given that we cannot find proof. As we discuss below, it is possible to entertain a contrary possibility, because powerful cues can bring "forgotten" information back into consciousness. Still, given the huge number of events in one's life, the idea that all would be stored forever (in some form) seems unlikely.

Forgetting as retrieval failure

Another possibility, essentially the obverse of the strong form of forgetting, might be considered a weak form of the concept. In its starkest form, this idea would maintain that all events that have been encoded and stored do somehow persist in the nervous system (including all those from age 7), and the inability to access them now is due to retrieval failure. Although this proposal might seem farfetched, when Loftus and Loftus (1980) surveyed psychologists many years ago, a large percentage (84%) favored something like this view. The percentage today might be lower, but the 1970s were the heyday of studies of retrieval in general and the power of retrieval cues in particular (Tulving & Thomson, 1973; for reviews see Roediger & Guynn, 1996; Tulving, 1983).

The idea of forgetting as retrieval failure is a scientifically useful concept, because (unlike the case with forgetting as storage failure) evidence can be found in its favor. Let us consider one experiment to demonstrate the point. Tulving and Pearlstone (1966) presented high-school students with lists of words to remember. Although there were many conditions, for our purposes consider the condition in which students studied 48 words that were members of 24 common categories, so they heard two words per category. Thus, students heard lists such as "articles of clothing: blouse, sweater; types of birds: blue jay, parakeet." The words were presented at a slow rate (2.5 sec/word) so the encoding of the words was ensured, in the sense that if the experimenter had stopped at any point, the subjects could have successfully recalled the last word presented. Thus, in this sense, all 48 words were learned.

One group of subjects was tested by free recall; they were given a blank sheet of paper and asked to recall the words in any order. They recalled 19.3

words, which means they forgot (failed to retrieve) about 29 others (28.7 to be exact). We can thus ask what happened to the forgotten words. It is logically possible that their representations had completely evaporated and had vanished from storage, but, as already discussed, we can never assume that. On the other hand, it could be that traces of the words were stored, but could not be retrieved with the minimal cues of free recall (people must use whatever cues they can internally generate). Tulving and Pearlstone (1966) found evidence for this latter possibility by giving the students (both the same group that had received a free recall test and a different group that had not had such a test) category names as cues. When the 24 category names (e.g., articles of clothing) were given, students were able to recall 35.9 words (and it did not matter much as to whether or not they had taken the prior free recall test). Thus, with stronger cues, students were able to recall nearly twice as many words as in free recall, showing that some of the forgetting in free recall was due to retrieval failures. Such powerful reversals of forgetting demonstrated in many experiments were probably why the psychologists surveyed in the late 1970s by the Loftuses claimed that forgetting was mostly due to retrieval failures.

Of course, even with the powerful category name cues, students still forgot about 25% of the words (12 of 48). Were these lost from storage? There is no way to know, but probably if the students had been further probed with recognition tests (with strong "copy cues") or with implicit tests (Schacter, 1987), evidence for storage of even more words would have been found. The asymmetry in the logic here – evidence of forgetting as retrieval failure can be obtained, but evidence of forgetting as storage failure cannot – leads back to Nietzsche's dictum. Still, as noted above, we cannot conclude that forgetting never involves elimination of stored traces, just that such a claim cannot be verified scientifically.

Forgetting as loss of information over time

A third way of defining forgetting, the one first used since Ebbinghaus (1885/1964) and many others since his time, is to plot retention of some experiences over time. This definition is complementary to the forgetting-as-retrieval-failure definition, not opposed to it. The typical way to conduct such forgetting experiments is to have (say) seven groups of subjects exposed to the same information (e.g., a list of words). One group would be tested immediately after learning, with other groups tested at varying delays after that point (e.g., 1 hour, 6 hours, 12, hours, 24 hours, 48 hours, and 1 week). Retention would be plotted across the various retention intervals and a forgetting curve would be derived, almost always showing less information recalled or recognized as a function of the time since learning. As Ebbinghaus (1885/1964) put it: "Left to itself every mental content gradually loses its capacity for being revived, or at least suffers loss in this regard under the influence of time" (p. 4). One critical methodological stricture in such experiments is that the type of test be held constant across delays, so that retrieval cues do not differ.

As noted, Ebbinghaus (1885/1964) was the first to plot forgetting over time. He presented his results in a series of tables in his book (see pp. 67–76), but later writers have chosen to show them in a figure and his findings appear in Figure 1.1. Ebbinghaus memorized lists of nonsense syllables so that he could recall them perfectly, and then he tried to relearn the list at varying delays from 19 minutes to 26 days. He measured the number of trials (or the amount of time) needed to learn the list perfectly in the first instance and then, later, he measured the trials or time to relearn the list after varying intervals. The measure shown in Figure 1.1 is percentage of savings in relearning the list, defined as the number of trials needed to learn the list originally (OL, for original learning) minus the number of trials needed for relearning (RL) divided by OL and then multiplied by 100 (to get a percentage). Thus, savings = (OL − RL)/OL × 100. Ebbinghaus noted that the shape of the forgetting curve appeared logarithmic.

This savings method of forgetting is not used much today, but nothing about the forgetting curve much hangs on the exact details of experimental design or the measure used, because nearly all forgetting functions look pretty much alike. Rubin and Wenzel (1996) examined "100 years of forgetting," seeking the best quantitative fit to the hundreds of forgetting curves that had been collected up until that point. They tried 105 different functions and concluded that 4 functions fit the forgetting curves quite well (and pretty much indistinguishably): the logarithmic function, the power function, the exponential in the square root of time, and the hyperbola in the square root of time. More recently, Wixted and Carpenter (2007) have argued that the power function is the correct one to describe the shape of the forgetting curve.

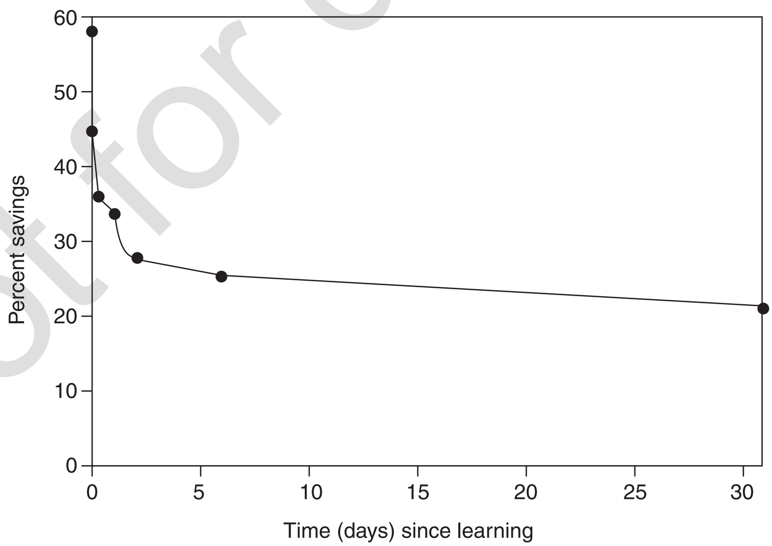

Figure 1.1 Forgetting curve adapted from Ebbinghaus (1885/1964, pp. 67–76).

Although curves like Figure 1.1 are called "forgetting curves," Rubin and Wenzel (1996) pointed out that this is a misnomer. These are retention curves, because the amount retained is plotted. If one really were to plot forgetting, then the curves would increase as a power function over time. True enough, but we will follow the common practice of calling such curves forgetting curves.

Most forgetting curves have been derived from verbal materials over periods of minutes to hours to days. However, even when radically different procedures are used, forgetting functions appear rather similar in that losses occur rapidly at first and then seem to approach an asymptote. The same shape occurs in loss of information from brief visual displays over a couple of seconds (Sperling, 1960), auditory presentations over about 4 seconds (Darwin, Turvey, & Crowder, 1972), holding a few items in short-term memory while distracted by another task (Peterson & Peterson, 1959), remembering a word over some minutes (Rubin, Hinton, & Wenzel, 1999), remembering lists over days (Slamecka & McElree, 1983), and remembering Spanish vocabulary learned in college over many years (Bahrick, 1984). Figure 1.2 shows data from the experiment by Rubin et al. (1999) just mentioned because they used ten measures to produce a more compelling curve than in many such experiments (often only a few data points are obtained).

Given the consistent forgetting effects shown in the literature, theories of forgetting have focused on the inexorable loss of information over time. We review below some of the main contending theories proposed to explain forgetting, but first we deal with a neglected side issue.

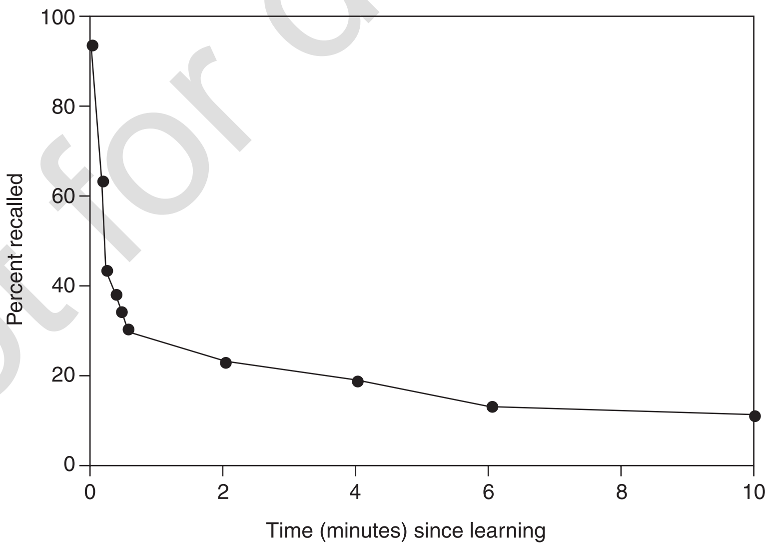

Figure 1.2 Forgetting curve adapted from Rubin, Hinton, and Wenzel (1999, Table A1, p. 1175).

A caveat

Ever since Ebbinghaus, forgetting experiments have employed one of two designs: either separate groups of subjects are exposed to the same material and tested at different points in time (a between-subjects design); or the same group of subjects is given many different sets of materials and the type of material tested at each delay is counterbalanced across subjects (a within-subjects, between-materials design). In both these cases, a particular set of materials is tested only once, because testing material may alter the forgetting curve. In fact, this concern is well founded, because testing does change the forgetting curve – tested material is subject to less forgetting than nontested material (e.g., Roediger & Karpicke, 2006). However, one might question whether the standard way of measuring forgetting, with people assessed only once on material, is particularly representative. After all, in life we all exist in "within-subject, within-materials" situations; for important memories, we recall them repeatedly; we repeatedly retrieve the events of our lives.

These considerations lead to the question of what happens when the same set of events (a list of words or pictures, or any other material) is repeatedly tested over time. Consider an experiment by Erdelyi and Becker (1974, Experiment 2) that meets the usual stricture of forgetting experiments: a set of material (either words or pictures) was presented to subjects and they were tested under the same conditions each time (with no cues provided). The only difference is that the subjects were tested three times, with each recall period occurring relatively soon after the prior recall period in one set of conditions in the experiment. The first test occurred shortly after study and lasted for 7 minutes. After that, the second test occurred for 7 more minutes, and then the third. Thus, as with customary forgetting studies, each successive test occurred after increasingly longer delays. The results are shown in Figure 1.3, where it can be seen that the "forgetting curves" look highly irregular. There was no forgetting of words, and recall of pictures actually improved across repeated tests at greater delays! Many others have replicated these results (e.g., Roediger & Thorpe, 1978) of increases in recall with repeated (and increasingly delayed) tests over time (see Payne, 1987; and Roediger & Challis, 1989, for early reviews of this literature, which actually dates back to early in the 20th century. Erdelyi, 1996 provides a more expansive review).

The pattern in Figure 1.3 indicates that, at the level of individual items, forgetting does not always occur over time because more items were recalled after longer intervals than shortly after learning. Thus, contrary to the quote from Ebbinghaus and much of the literature on forgetting curves, at the level of individual items there is no inexorable decline in "trace strength" or else an item could not be recovered at a later time that was not recalled at an earlier time. This claim is obviously true in the case of pictures from the data in Figure 1.3, but it turns out to be true (at the item level, if not always the list level) in the case of words, too. That is, on a second test, both individual

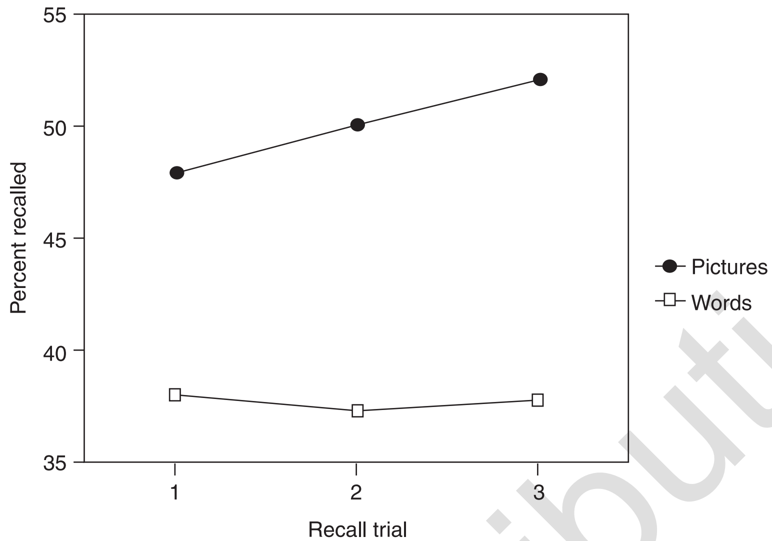

Figure 1.3 Data adapted from Erdelyi and Becker (1974, Figure 1, no interval group, p. 165).

words and pictures that were not recalled on the first test can be recalled on the second test (this phenomenon is called reminiscence; Roediger & Thorpe, 1978). Yet items are also forgotten between tests, and in the case of recall of words in Erdelyi and Becker's (1974) experiment, these two quantities (intertest forgetting and intertest recovery or reminiscence) offset one another for no net increase or decrease. However, in the case of pictures, recovery of items between tests was greater than forgetting, so a net increase occurred. Erdelyi and Becker (1974) labeled this net increase hypermnesia, the improvement in recall over time with repeated tests. Others have reported hypermnesia for words and other sorts of material, too (see Payne, 1987).

Much research has been conducted on the topic of reminiscence and hypermnesia, but this literature has not been incorporated into the study of forgetting for the very good reason that it does not fit. Most writers do not even consider it, but Underwood (1966) at least noted its existence in his chapter on forgetting in his popular textbook. He then went on to say: "We will not be concerned with reminiscence in this chapter" (p. 544), which is one way to deal with the problem (even though not a particularly satisfactory one). Still, it is understandable, because theories of forgetting are mute about improvements in performance with delays from initial learning. Other traditions of work showing such improvements over time exist, too – spontaneous recovery in animal and human learning, reminiscence in motor learning, enhanced performance in motor skill learning after sleep, among others. Wheeler (1995) provided some review and evidence for spontaneous recovery in an interference paradigm.

Although researchers studying forgetting ignore the hypermnesia literature – none of the other authors in this volume touch on the issue – we believe it should be considered. The very facts of reminiscence and hypermnesia point to the importance of retrieval factors and support the definition of forgetting as retrieval failure (Tulving & Pearlstone, 1966). One basic idea is that of a limited capacity retrieval system (Tulving, 1967) in which we know (have stored) much more than we can retrieve at any point in time. Retrieval forms a bottleneck in the system, a fundamental limitation. As discussed below, retrieval of some information often causes forgetting of other information, so that retrieval becomes a self-limiting process (Roediger, 1974, 1978; see too Bjork, Bjork, & Caughey, 2007).

Theories of forgetting: A brief tour

This entire book is about theories of forgetting. Here we set the stage by discussing, quite briefly, the main theories.

Decay theory

This is the oldest and simplest theory, which states that forgetting occurs because of the "wasting effects of time" (McGeoch, 1932). This theory essentially amounts to saying that "forgetting happens." The analogy sometimes made is that memories are like muscles and they atrophy (decay) if they are not used, so they grow ever weaker over time, although this statement merely describes the forgetting curve without explaining why it occurs.

In a classic paper, McGeoch (1932) mounted a withering attack on decay theory from which it has never really recovered. First, he argued that it was improper as a scientific theory because it did not specify a mechanism by which the memory trace would unwind over time. Second, he pointed to data from experiments showing reminiscence (e.g., Brown, 1923) in which items not recalled at one point in time could be recalled later, which is completely inconsistent with decay theory. (This is the point raised in the previous section.) And third, he argued that even when passage of time was controlled, forgetting could be determined by the number or density of events during that time; the more events, the greater forgetting. He pointed to Jenkins and Dallenbach's (1924) experiments showing that greater forgetting of verbal materials occurred after equivalent periods of waking than of sleep. These data are shown in Figure 1.4 (the data were obtained from Dallenbach, 1963). The effects of sleep on retention are a topic of lively interest on the contemporary scene and are discussed in detail by Peigneux, Schmitz, and Urbain in this volume (Chapter 8). Brown and Lewandowsky (Chapter 4) hammer another nail or two into decay theory's coffin.

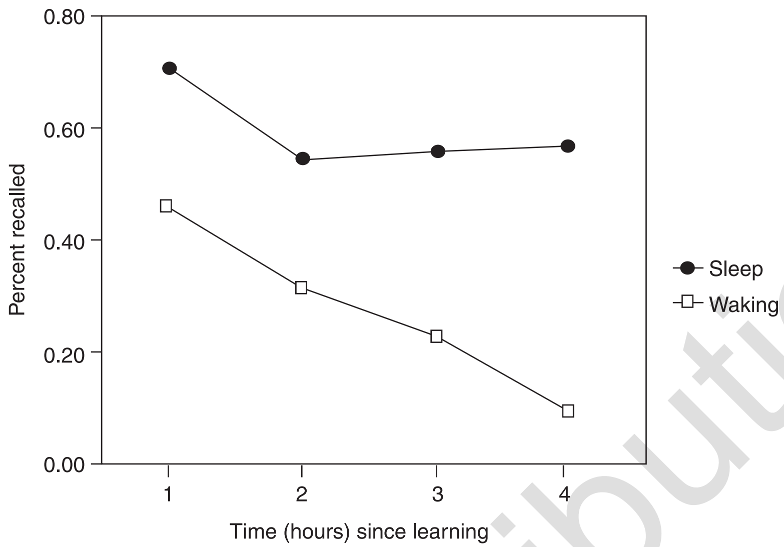

Figure 1.4 Data adapted from Dallenbach (1963, Table 1, p. 701).

Interference theory

While McGeoch (1932) was torching decay theory, he argued that the most important factor in forgetting was actually interference. Interference can arise from anything other than the to-be-remembered information. McGeoch (1942) put it in a straightforward manner in saying that forgetting is often a result of the wrong memory being accessed by a particular cue.

Interference can take many forms, but has been broadly divided into two types: proactive and retroactive. *Proactive interference* refers to the negative effects of prior learning on retention of target information, whereas *retroactive interference* refers to the negative effects of encountering new information after encoding target information. If you drive to work and park in various spots in the same parking lot every day, imagine someone asking you in what spot you parked one week ago. Even if you found your car perfectly that day (thus indicating that you had encoded and stored the information well enough to retrieve it hours later), you would probably have trouble recalling, a week later, where your car was parked on that day. According to interference theorists, the forgetting is due to two sources: all the times you parked in the lot before the critical date create proactive interference, whereas your comings and goings of the past week provide retroactive interference. McGeoch (1932, 1942) argued that retroactive interference was the most potent cause of forgetting.

Retroactive interference

The simplest way to demonstrate retroactive interference is to get subjects to learn a cue–target association (for example, horse–umbrella, or A–B), with the eventual test being to recall B (umbrella) when given A (horse). Two groups of subjects learn a list of A–B pairs to one perfect recitation. Then, in an experimental condition, subjects learn conflicting pairs (A–D, like horse–automobile), again so that they know the pairs perfectly. New responses are paired with the same cues. In a different (control) group subjects learn new pairs after the original A–B learning, so they might learn piano–automobile (C–D learning, where C–D is unrelated to A–B). After both groups have learned their second lists, a delay occurs. The final, criterial task is for both groups to receive the original A cues (horse) with instructions to recall items from the first list. Their task at test is to recall the targets that were paired with the cues in list 1. The finding is that subjects who have experienced the A–B, A–D arrangement recall the responses from the original list less well than those in the A–B, C–D condition. This outcome defines retroactive interference.

Another control condition is sometimes used in which either no activity or a general distracter task (e.g., reading a book or playing a videogame) is employed after A–B learning. Usually this condition produces little forgetting of the A–B pair. The general finding is that, relative to no activity, learning C–D pairs after A–B learning decreases probability of recall of B somewhat, but A–D learning causes much more forgetting. The former type of forgetting is referred to as “nonspecific interference,” whereas the latter is caused by “specific interference” (because the A–D pair specifically conflicts with A–B recall).

Two primary processes have been used to explain retroactive interference: unlearning and response competition (Melton & Irwin, 1940). These two comprise the processes of “two factor interference theory.” The basic idea for unlearning is that the A–B association is weakened or destroyed when A–D is learned (reminiscent of Nietzsche’s definition of forgetting). However, a different view is that the A–B association remains as A–D is learned, but the responses compete with one another during retrieval in response to the cue A–???. This factor endorses the idea of forgetting as retrieval failure.

From the 1940s through the 1960s, researchers used paired associate paradigms to seek evidence for the two factors thought to be responsible for retroactive interference. Crowder (1976, Chapter 8) provides a thorough history of the work through the early 1970s, and Wixted (this volume, Chapter 13) helps to bring the discussion to the present.

Proactive interference

From 1932 through the mid-1950s, proactive interference received short shrift in discussions of forgetting. The discovery that powerful effects of interference from events occurring *prior* to learning some target events can be attributed to Underwood (1957), who set out to solve the conundrum of why subjects in various studies showed remarkably different rates of forgetting over a 24-hour period. He demonstrated that by far the largest factor in forgetting of word lists over a day was the number of word lists studied *before* rather than *after* the target list.

This outcome can be demonstrated using the paradigm described above (often referred to as the A–B, A–D paradigm), whereby the set-up is exactly the same, except that subjects are now tested on list 2 (i.e., they are asked to produce target D when given A) after 24 hours. Underwood (1957) was puzzled by the fact that forgetting in this design differed so dramatically from study to study. However, after some careful scientific detective work (a kind of early meta-analysis), he discovered that the critical variable was the number of prior lists that subjects had learned before the critical list on which they were to be tested the next day. In reviewing the literature, he found that when subjects had learned 15–20 lists prior to learning a last list perfectly, they recalled only 15–20% of the list a day later. However, if subjects learned only one list on the first day, they recalled 80–85% after 24 hours. Of course, according to two-factor theory, proactive interference must be due to response competition, because unlearning does not apply in the proactive case.

Following Underwood's (1957) report, proactive interference became much more studied. However, findings such as those from Jenkins and Dallenbach's (1924) sleep study (described above) and many more studies showed that retroactive interference was still a critical factor in forgetting. In addition, Underwood and Postman (1960) launched a theory arguing that proactive interference from prior linguistic habits was critical to forgetting in laboratory paradigms, but they were later forced to abandon this theory in response to negative evidence (see Crowder, 1976, for a good account of this story).

Wixted (2004) has proposed that the field's concentration on proactive interference was a mistake that possibly led to the "demise" of interference theories of forgetting. He has gone so far as to argue that the whole A–B, A–D list learning paradigm and the tradition surrounding it "may pertain mainly to forgetting in the laboratory and that everyday forgetting is attributable to an altogether different kind of interference" (p. 235). Wixted (Chapter 13) revisits these historical developments and suggests a new role for interference in forgetting that takes into account recent psychological and neuroscientific developments. This chapter can be read alongside Brown and Lewandowsky (Chapter 4) who take a different position. In addition, Dewar, Cowan, and Della Sala (Chapter 9) apply the concept of retroactive interference to explaining anterograde amnesia.

Wixted (2004) may have been hasty in dismissing classic interference theories as irrelevant to forgetting outside the laboratory. It is useful to consider an earlier example of these ideas being written off. In an influential article from several decades ago, Neisser (1978) castigated both learning theory in general and interference theory in particular by saying: "With learning theory out of fashion, the experiments of the interference theorists seem like empty exercises to most of us. Were they ever anything else?" (from p. 8 of Neisser, 1982, which reprinted the 1978 chapter).

Neisser (1978) was arguing that laboratory approaches to studying human memory should be avoided in favor of naturalistic (or at least more realistic) studies. In the same essay (p. 15), he extolled the virtues of Loftus and Palmer's (1974) interesting studies of eyewitness memory (and forgetting). In retrospect, this juxtaposition seems ironic, because the Loftus tradition of studying eyewitness memory actually depends on similar processes to those in classic studies of retroactive interference. In Loftus and Palmer's classic misinformation experiments, which have been repeated in various ways many times, subjects are presented with slides that tell a story about, say, a traffic accident. Following the slide presentation, they are exposed to a test or a passage that contains some inconsistencies with the original story presented in the slides – in other words, misinformation. For instance, in the slide show subjects may have seen a picture of a car driving by a STOP sign, while the text read later refers to a YIELD sign instead. On a final test, subjects are given a choice between the two types of signs and asked to indicate which one they saw in the slides (or they may be asked to recall the type of sign). The outcome is that, relative to a control condition in which the sign was referred to in some neutral way ("a traffic sign"), subjects given the misinformation are much more likely to falsely remember the sign as a YIELD sign (in this example). The misinformation leads to errors in the witness's memory, which has obvious implications for eyewitness testimony in court.

A critical issue is why such errors occur in eyewitness memory: What happens to the original memory for the STOP sign in the slides when subjects incorrectly remember the YIELD sign as a result of the misinformation? Has this memory been inexorably forgotten (although we may never be able to prove it), has it been somehow altered, or is it intact but temporarily inaccessible due to competition from the YIELD sign? These controversies have exact parallels in the retroactive interference literature (see Roediger, 1996). After all, the Loftus paradigm can be considered a species of retroactive interference of the A–B, A–D variety: study sign–STOP, study sign–YIELD, then recall (or recognize) the first kind of sign on a later test.

Loftus and her colleagues originally interpreted their results as showing that the original trace had been changed (from a representation of a stop sign to a yield sign), which is akin to the unlearning interpretation of retroactive interference. However, McCloskey and Zaragoza (1985) later argued that nothing had happened to the original trace, but the interference that occurred in the Loftus misinformation paradigm came about because of competition between responses (STOP and YIELD), the second factor in classic two-factor interference theory. The debate in the misinformation paradigm over the years has recapitulated in many ways the arguments from classic studies of interference from the 1950s and 1960s (Roediger, 1996).

Input and output interference

Although proactive and retroactive interference are well known as possible causes of forgetting, Tulving and Arbuckle (1963) pointed to two other (complementary) sources of forgetting: input and output interference. They discussed these as sources of interference within a single trial. Item 3 in a list of 5 items will be better recalled than item 3 in a list of 10 items; the more events occurring, the less the probability of recalling any one event, which is the operational definition of input interference. Input interference refers to the fact that for larger sets of to-be-learned material, the greater the probability of forgetting any particular item in the set (all other things being equal). This observation forms part of the basis for cue overload theory, discussed below.

The concept of output interference has perhaps enjoyed a more exciting fate as a cause of forgetting than input interference, albeit in a somewhat different incarnation than Tulving and Arbuckle (1963) originally envisioned. Their original idea was that the more items tested before any particular item (the more items "output"), the worse would recall be for the next item. Their experimental situation involved short-term recall, so the act of recall could be considered as a distracter task that eliminated information from primary (or short-term or working) memory. However, the same idea operates in long-term memory (e.g., Brown, 1968; Roediger, 1974) and is now often called retrieval-induced forgetting, due to the influential experimental and theoretical work of Anderson, Bjork, and Bjork (1994) and Anderson and Spellman (1995).

In the retrieval-induced forgetting paradigm introduced by Anderson et al. (1994), subjects were presented with word pairs consisting of category names and exemplars of the category (e.g., furniture–chair; furniture–table; fruit–banana; fruit–apple). Following initial study, they were then given a chance to practice some items from these categories, but only certain exemplars were practiced. In our example, they might practice the furniture category, but they would be repeatedly cued with items like furniture–c______ and retrieve chair. However, other category members (table, in our example) would remain unpracticed. On a later test subjects were given category names and asked to recall all items from the category. The finding is that the items from the practiced category show two effects relative to retrieval from the unpracticed category (from the fruit category, in our example, where no items were practiced). First, the previously practiced items are recalled better than those from the nonpracticed category, in line with work discussed earlier on the effects of testing on retrieval (Roediger & Karpicke, 2006). Second, and more importantly for present purposes, nonpracticed items in the practiced category (like table in our example) were recalled more poorly than items from the nonpracticed category (the fruit items in our example). Thus, active retrieval of some items from the category induced forgetting of the other items, hence the name of the phenomenon: retrieval-induced forgetting.

This ability of our memories to actively inhibit information is crucial for avoiding cognitive overload and producing appropriate responses to the environment. Harris, Sutton, and Barnier (Chapter 12) explore individual differences in retrieval-induced forgetting and how this phenomenon may map onto autobiographical memory.

Retrieval theories

In McGeoch's famous 1932 paper on forgetting, he mentioned (almost in passing) that he believed "altered stimulating conditions" between the context of learning and that of use of information (retrieval) was a cause of forgetting (in addition to retroactive interference). He meant that retention would be better the more the conditions at test matched those during learning and that, conversely, changed conditions between learning and testing would lead to forgetting. In the 1970s, this basic idea attained new adherents as Tulving and Thomson (1973) proposed the encoding specificity principle as governing the effectiveness of retrieval cues. The basic claims are that events are encoded in terms of specific patterns of features; that cues in the retrieval environment also are encoded as feature bundles at the time of retrieval; and (critically) to the extent that features in the cues overlap or match those in the trace, memories for experiences will be evoked (see Flexser & Tulving, 1978, for a formal instantiation of these ideas). These ideas formalize McGeoch's offhand comment and are critical for retrieval analyses of forgetting, which tacitly assume availability of trace information that must be matched by information in cues for retrieval to occur (Tulving, 1983). Forgetting over time may be due to loss of information in the trace or to increasing mismatch between cues and the information in the trace, according to retrieval theories. Much evidence supports these basic ideas (Roediger & Guynn, 1996; Tulving, 1983).

Cue overload

Another theory of forgetting, complementary to interference theory and emphasizing retrieval factors, is cue overload theory (Earhard, 1967; Watkins & Watkins, 1975). The basic idea is straightforward: the more events that are subsumed under a particular cue, the greater the likelihood of forgetting an item associated with a cue. For example, in the A–B, A–D paradigm, two target events are attached to the same cue and hence each is less memorable than if only one were attached. In a different situation, if a list contains many types of furniture, the retrieval cue "furniture" will be less effective at provoking recall of any particular instance of furniture than if a list had presented only one or two types of furniture (e.g., Roediger, 1973). Watkins (1979) provided further examples of this principle in action. The use of the cue overload principle has become ubiquitous in research on forgetting and especially to interference paradigms (see Wixted, Chapter 13). It is useful, if descriptive.

Consolidation

A critical concept in the science of memory is consolidation (e.g., Nadel, 2008). Consolidation may be defined as “the progressive post-acquisition stabilization of the engram” and/or “the memory phase(s) during which [it] takes place” (Dudai, 2002, p. 59). Thus, forgetting may occur because engrams or memory traces are labile; they may last briefly and support retention over the short term, but unless consolidation occurs, the memories will be forgotten.

Related to the issue of forgetting is research on the molecular process of reconsolidation (Sara, 2000, 2008). The idea is that each time a memory is retrieved, it undergoes the same sort of molecular process that happens after initial encoding. Crucially, if this process is interfered with (which can be done by means of chemical inhibitors, see Dudai, 2006), the memory can become altered or, hypothetically, even lost. This basic idea should seem familiar, because we have met it in A–B, A–D interference studies and in the Loftus misinformation work; events coming after a target event may somehow undo or interfere with the target memory. Many chapters in the current volume expand on ideas of lack of consolidation and/or reconsolidation as causes of forgetting (e.g., Wixted, Chapter 13).

Repression

The concept of repression is used to explain some types of forgetting. Freud (1914/1957) popularized the idea that forgetting may be motivated by a need to protect the psyche from threatening memories or thoughts. The idea predates Freud, but he brought it into prominence and adduced many clinical case studies that he thought supported the concept. However, it has had a controversial history. To complicate matters, repression may be defined in several different ways, and Freud changed his theoretical ideas several times during the course of his long career.

At the simplest level, repression is the process of trying to avoid painful memories. So, if a person has bad experiences at work one day and decides to watch a lighthearted movie that evening to put aside (to forget about) the events of the day, that activity would meet this very weak definition of repression. If this were all that were meant by repression, it would not be controversial. Similarly, motivated forgetting of the sort of failing to remember a dentist appointment and thus avoiding pain would fall into this garden variety example of repression.

A second definition of repression holds that ideas and memories may be firmly held in a conscious state, then banished from consciousness into an unconscious state and hence forgotten. This suppression is an active, effortful process, but once the memories become unconscious, they reveal themselves only indirectly (e.g., through Freudian slips or through dreams, the “royal road to the unconscious”). Unconscious memories can also cause unwanted effects on experience and behavior and thus be the source of various mental and even physical problems. We will not discuss the issue of repression as a cause of forgetting further in this context (none of the authors of this volume addresses the idea), but historically the idea of repression has played a central role in certain aspects of psychological theorizing and experimentation (see Erdelyi, 1985).

Adaptive reasons for forgetting

So much angst has been expressed about the erroneous nature of human memory – both in terms of forgetting and, perhaps even worse, the creation of false memories – that we might wonder why our memories have evolved to be so fragile and fallible. However, once we pause to consider the adaptive nature of forgetting and interference, we can see plausible reasons that forgetting exists. For example, if we move to a new city, we must learn a new address and telephone number (among many other things) and not have the old ones constantly intruding. We need to forget them, even though they are well learned. More generally, as our environment changes, so must our memories. People who cannot forget are often plagued with problems, as in Luria's (1968) classic study of S, a mnemonist whose synesthesia empowered (or overpowered) him with a strikingly good recollection of even trivial events from his life. More recently, Parker, Cahill, and McGaugh (2006) reported the case of a woman plagued by the inability to forget the happenings of her life.

In order to understand the value of forgetting, we need to take a step back and consider the function of memory outside the context of attempts to remember autobiographical events or word lists in an experiment. It is likely that our capacity to remember evolved as a tool for navigating the present and planning for the future, rather than for looking back on the past (Nairne & Pandeirada, 2008). To this end, it is not practical or useful to maintain detailed, veridical information of encoded events and information in memory indefinitely. Instead, Bjork and Bjork (1988) have proposed that "disused" memories – those that are retrieved less and less over time, such as the address of your childhood home – become less accessible in order to allow for more relevant information, such as your current address, to take precedence. Crucially, the loss of access to information through disuse is seen not as a failure of the system, but an adaptive feature that facilitates updating (Bjork, 1978).

Anderson and Schooler (1991) provided a more formal analysis of the adaptive nature of forgetting by demonstrating striking parallels between the statistical occurrence of events in the environment and the typical negatively accelerated retention function shown in Figure 1.1. The idea is that events that have been occurring frequently in the recent past are also more likely to occur in the near future. For instance, Anderson analyzed his own email inbox and discovered that on a given day he was more likely to receive an email from someone who had written him recently (and generally more often in the recent past) than someone who had only written a while back. Hence, at any given moment, he was more likely to require access to information about recent senders. The same was true of many other sets of data that Anderson and Schooler examined. While the mathematical analyses involved in Anderson and Schooler's theory are far beyond the scope of this chapter, the take-home message is that forgetting may not be an accident of nature. Rather, the forgetting function may be shaped to mirror the frequency of events in the environment and how they change over time.

Although not usually considered in evolutionary terms, many laboratory phenomena may reveal positive adaptations of forgetting. Retroactive interference can be considered an adaptation if old (unneeded) information is replaced by new, updated information, as in the examples of learning new addresses and telephone numbers.

Conclusion

The aim of this chapter was to provide a brief overview of some key issues in the scientific study of forgetting, but it is by no means complete. Our chapter has focused primarily on the experimental psychologist's approaches to studying forgetting, but, as the remaining chapters in this volume indicate, numerous approaches exist. Although we discussed consolidation and reconsolidation rather tersely, these issues occupy many writers in this book. In fact, as many as six chapters (Brown & Lewandowsky, Chapter 4; Murre, Chapter 5; Levy, Kuhl, & Wagner, Chapter 7; Peigneux et al., Chapter 8; Dewar et al., Chapter 9; and Wixted, Chapter 13) deal extensively with the issue of consolidation in relation to domains ranging from sleep (Chapter 8) to amnesia (Chapter 9). Although we have focused primarily on behavioral data from healthy adults, other chapters in this volume present new and fascinating perspectives on forgetting in patients with Alzheimer's disease (Dewar et al., Chapter 9), epilepsy (Butler, Muhlert, & Zeman, Chapter 10), and psychogenic amnesia (Brand & Markowitsch, Chapter 11), as well as forgetting theories based on alternative techniques including connectionist modeling (Murre, Chapter 5) and neuroimaging (Levy, Kuhl, & Wagner, Chapter 7). Despite the fact that our chapter is incomplete, the issues revolving around the definition and leading theories of forgetting must be borne in mind for all treatments of the topic.

References

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, *2*, 396–408.

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

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

Bahrick, H. P. (1984). Semantic memory content in permastore: Fifty years of memory for Spanish learned in school. *Journal of Experimental Psychology: General*, *113*, 1–29.

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 II* (pp. 283–288). Chichester: 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., Bjork, E. L., & Caughey, J. B. (2007). Retrieval as a self-limiting process: Part II. In J. S. Nairne (Ed.), *The foundations of remembering: Essays in honor of Henry L. Roediger, III* (pp. 19–37). New York: Psychology Press.

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

Brown, W. (1923). To what extent is memory measured by a single recall? *Journal of Experimental Psychology*, *6*, 377–382.

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

Dallenbach, K. M. (1963). Tables vs. graphs as means of presenting experimental results. *American Journal of Psychology*, *76*, 700–702.

Darwin, C. J., Turvey, M. T., & Crowder, R. G. (1972). An auditory analogue of the Sperling partial report procedure: Evidence for brief auditory storage. *Cognitive Psychology*, *3*, 255–267.

Davis, M. (2008). Forgetting: Once again, it's all about representations. In H. L. Roediger, Y. Dudai, & S. M. Fitzpatrick (Eds.), *Science of memory: Concepts* (pp. 317–320). New York: Oxford University Press.

Dudai, Y. (2002). *Memory from A to Z: Keywords, concepts, and beyond*. Oxford: Oxford University Press.

Dudai, Y. (2006). Reconsolidation: The advantage of being refocused. *Current Opinion in Neurobiology*, *16*, 174–178.

Earhard, M. (1967). Cued recall and free recall as a function of the number of items per cue. *Journal of Verbal Learning and Verbal Behavior*, *6*, 257–263.

Ebbinghaus, H. (1885/1964). *Memory: A contribution to experimental psychology*. Oxford: Dover.

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

Erdelyi, M. H. (1996). *The recovery of unconscious memories: Hypermnesia and reminiscence.* Chicago: University of Chicago Press.

Erdelyi, M. H., & Becker, J. (1974). Hypermnesia for pictures: Incremental memory for pictures but not words in multiple recall trials. *Cognitive Psychology*, *6*, 159–171.

Flexser, A. J., & Tulving, E. (1978). Retrieval independence in recognition and recall. *Psychological Review*, *85*, 153–171.

Freud, S. (1914/1957). The history of the psychoanalytic movement. In J. Strachey (Ed.), *The standard edition of the complete psychological works of Sigmund Freud* (Vol. 14). London: Hogarth Press.

Jenkins, J. G., & Dallenbach, K. M. (1924). Obliviscence during sleep and waking. *American Journal of Psychology*, *35*, 605–612.

Köhler, W. (1947). *Gestalt psychology: An introduction to new concepts in modern psychology*. New York: Liveright.

Loftus, E. F., & Loftus, G. R. (1980). On the permanence of stored information in the human brain. *American Psychologist*, *35*, 409–420.
Loftus, E. F., & Palmer, J. C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior*, *13*, 585–589.
Luria, A. R. (1968). *The mind of a mnemonist: A little book about a vast memory*. Cambridge, MA: Harvard University Press.
McCloskey, M., & Zaragoza, M. (1985). Misleading postevent information and memory for events: Arguments and evidence against memory impairment hypotheses. *Journal of Experimental Psychology: General*, *114*, 1–16.
McGeoch, J. A. (1932). Forgetting and the law of disuse. *Psychological Review*, *39*, 352–370.
McGeoch, J. A. (1942). *The psychology of human learning*. Oxford: Longmans, Green, and Company.
Melton, A. W. (1963). Implications of short-term memory for a general theory of memory. *Journal of Verbal Learning and Verbal Behavior*, *2*, 1–21.
Melton, A. W., & Irwin, J. M. (1940). The influence of degree of interpolated learning on retroactive inhibition and the transfer of specific responses. *American Journal of Psychology*, *53*, 173–203.
Nadel, L. (2008). Consolidation: The demise of the fixed trace. In H. L. Roediger, Y. Dudai, & S. M. Fitzpatrick (Eds.), *Science of memory: Concepts* (pp. 177–181). New York: Oxford University Press.
Nairne, J. S., & Pandeirada, J. N. S. (2008). Forgetting. In H. L. Roediger (Ed.), *Learning and memory: A comprehensive reference* (Vol. 2, pp. 179–194). Oxford: Elsevier.
Neisser, U. (1978). Memory: What are the important questions? In M. M. Gruneberg, P. E. Morris, & R. N. Sykes (Eds.), *Practical aspects of memory* (pp. 3–24). London: Academic Press.
Neisser, U. (1982). *Memory observed: Remembering in natural contexts*. San Francisco: W. H. Freeman.
Parker, E. S., Cahill, L., & McGaugh, J. L. (2006). A case of unusual autobiographical remembering. *Neurocase*, *12*, 35–49.
Payne, D. G. (1987). Hypermnesia and reminiscence in recall: A historical and empirical review. *Psychological Bulletin*, *101*, 5–27.
Peterson, L. R., & Peterson, M. J. (1959). Short-term retention of individual verbal items. *Journal of Experimental Psychology*, *58*, 193–198.
Roediger, H. L. (1973). Inhibition in recall from cueing with recall targets. *Journal of Verbal Learning and Verbal Behavior*, *12*, 261–269.
Roediger, H. L. (1974). Inhibiting effects of recall. *Memory & Cognition*, *2*, 261–269.
Roediger, H. L. (1978). Recall as a self-limiting process. *Memory & Cognition*, *6*, 54–63.
Roediger, H. L. (1996). Memory illusions. *Journal of Memory and Language*, *35*, 76–100.
Roediger, H. L., & Challis, B. H. (1989). Hypermnesia: Increased recall with repeated tests. In C. Izawa (Ed.), *Current issues in cognitive processes: The Tulane Flowerree symposium on cognition* (pp. 175–199). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
Roediger, H. L., & Guynn, M. J. (1996). Retrieval processes. In E. L. Bjork & R. A. Bjork (Eds.), *Human memory* (pp. 197–236). San Diego: Academic Press.

Roediger, H. L., & Karpicke, J. D. (2006). The power of testing memory: Basic research and implications for educational practice. *Perspectives on Psychological Science*, *1*, 181–210.

Roediger, H. L., & Thorpe, L. A. (1978). The role of recall time in producing hypermnesia. *Memory & Cognition*, *6*, 296–305.

Rubin, D. C., Hinton, S., & Wenzel, A. (1999). The precise time course of retention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 1161–1176.

Rubin, D. C., & Wenzel, A. E. (1996). One hundred years of forgetting: A quantitative description of retention. *Psychological Review*, *103*, 734–760.

Sara, S. J. (2000). Retrieval and reconsolidation: Toward a neurobiology of remembering. *Learning & Memory*, *7*, 73–84.

Sara, S. J. (2008). Consolidation: From hypothesis to paradigm to concept. In H. L. Roediger, Y. Dudai, & S. M. Fitzpatrick (Eds.), *Science of memory: Concepts* (pp. 183–189). New York: Oxford University Press.

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

Sills, D. L., & Merton, R. K. (1968). *International encyclopedia of the social sciences*. New York: Macmillan Free Press.

Slamecka, N. J., & McElree, B. (1983). Normal forgetting of verbal lists as a function of their degree of learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *9*, 384–397.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, *74*, 1–30.

Tulving, E. (1967). The effects of presentation and recall of material in free-recall learning. *Journal of Verbal Learning and Verbal Behavior*, *6*, 175–184.

Tulving, E. (1974). Cue-dependent forgetting. *American Scientist*, *62*, 74–82.

Tulving, E. (1983). *Elements of episodic memory*. New York: Oxford University Press.

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

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

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

Underwood, B. J. (1957). Interference and forgetting. *Psychological Review*, *64*, 49–60.

Underwood, B. J. (1966). *Experimental psychology*. New York: Appleton-Century-Crofts.

Underwood, B. J., & Postman, L. (1960). Extra-experimental sources of interference in forgetting. *Psychological Review*, *67*, 73–95.

Watkins, M. J. (1979). Engrams as cuegrams and forgetting as cue overload: A cueing approach to the structure of memory. In C. R. Puff (Ed.), *Memory organization and structure* (pp. 347–372). New York: Academic Press.

Watkins, O. C., & Watkins, M. J. (1975). Buildup of proactive inhibition as a cue-overload effect. *Journal of Experimental Psychology: Human Learning and Memory*, *1*, 442–452.

Weiner, B. (1966). Effects of motivation on the availability of memory traces. *Psychological Bulletin*, *65*, 24–37.

Wheeler, M. A. (1995). Improvement in recall over time without repeated testing:

Spontaneous recovery revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 173–184.

Wixted, J. T. (2004). The psychology and neuroscience of forgetting. *Annual Review of Psychology*, *55*, 235–269.

Wixted, J. T., & Carpenter, S. K. (2007). The Wickelgren power law and the Ebbinghaus savings function. *Psychological Science*, *18*, 133–134.