2.12 Forgetting

J. S. Nairne and J. N. S. Pandeirada, Purdue University, West Lafayette, IN, USA

2.12.1 Introduction

As Ebbinghaus observed in his famous monograph, "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" (1885/1964: 4). To most, forgetting is a scourge, a nuisance, a breakdown in an otherwise efficient mental capacity. The momentary loss of information is a regular part of the human experience, but normal forgetting, even when permanent, is misconstrued as a malfunction or breakdown. We forget for adaptive reasons, and understanding the characteristics of forgetting ultimately tells us a great deal about why and how we remember.

In this chapter we offer a general tutorial on the psychology of forgetting. Our focus will be on its empirical characteristics and proposed theoretical underpinnings, as revealed primarily through laboratory studies with healthy human participants. Excellent reviews on abnormal forms of forgetting, such as those that occur in brain-damaged patients, and on forgetting in nonhuman populations can be found elsewhere in this series. We begin the chapter with a brief discussion of the adaptive value of forgetting, followed by an examination of its functional and mathematical characteristics. Next, we discuss possible causal mechanisms in some detail – why do we forget? Finally, we end by reconsidering the meaning of forgetting and its proper role in modern memory theory.

2.12.2 Forgetting and Its Adaptive Value

At face value, memory seems to be about recovering the past, recapturing or reviving previous experiences. Yet, it is unlikely that memory actually evolved for this specific purpose. The past can never occur again, at least in exactly the same form, so there is limited adaptive value in developing a system that carries around intact records of prior experiences. Instead, memory has value because it allows us to use the past in the service of the present, to decide on an appropriate plan of action now or in the future (Suddendorf and Corballis, 1997; Tulving, 2002). Intact records of the past are relevant in some situations, but not in others. For example, we might need to remember the specific location of a food source, but we need not remember every instance in which a particular food type was consumed (we need to remember only that it was edible). As Lewis Carroll famously quipped, "it's a poor sort of memory that only works backwards."

Recognizing that memory's primary function is to deal with the present, or perhaps to anticipate the future, informs how we need to think about forgetting. Obviously, if nature did not 'design' memory to reproduce the literal past, then it is not surprising that we sometimes have difficulty reproducing it. The veridical details of an event tend to be ignored by our memory systems, which choose instead to process and store inferences or connections that are likely to benefit future responding (Bartlett, 1932; Schacter and Addis, 2007). Even when the details of an event are correctly stored, there is little reason to anticipate that those details will be stored indefinitely. In fact, it is easy to make the case that forgetting is a highly adaptive feature of cognition (Bjork and Bjork, 1988, 1996; Altmann and Gray, 2002). Having an intermediate, rather than complete, retention of the past may improve our ability to use inferential heuristics (Schooler and Hertwig, 2005), maximize our ability to detect causality (Kareev, 2000), and even maintain a sense of sanity in an ever-changing world (Luria, 1968).

More to the point, a well-designed memory system can be expected to show sensitivity to the likelihood that a past event will be needed, or appropriate, to a future situation. It makes no sense to retain long-defunct telephone numbers, or high school locker combinations: These are more apt to produce needless clutter than potentially useful records. Moreover, once an event occurs in the present, the likelihood that it will occur again (at least in a similar form) changes predictably with time. If a predator appears in your environment at time t, then the chances that it will appear again are usually greater at time $t+1$ than at time $t+2$. As it turns out, the function relating event recurrence with time typically takes a negatively accelerated form, just like the classic forgetting function (Anderson and Schooler, 1991). What we normally think of as forgetting, therefore, may simply represent memory's sensitivity to the statistical structure of events in the environment. We forget an item's occurrence with time because, in fact, that item is less likely to occur again with time.

2.12.3 The Characteristics of Forgetting

The fact that forgetting takes on a characteristic form is *prima facie* evidence that a psychological process is at work. We do not forget things randomly; rather, the loss of information proceeds in an understandable and predictable fashion. A typical retention function is shown in **Figure 1** and it contains a stark empirical regularity: There is a negatively accelerating downward loss in retention over time. We forget rapidly at first and then retention slowly levels off. This general pattern occurs regardless of the quantity or quality of the information learned and the retention measure employed.

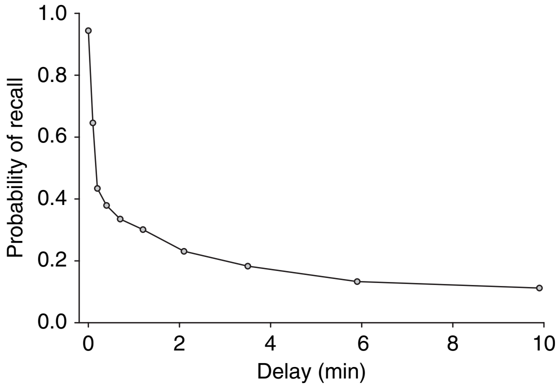

Figure 1 Probability of correctly recalling a word as a function of time. Data from Rubin DC, Hinton S, and Wenzel AE (1999) The precise time course of retention. *J. Exp. Psychol. Learn. Mem. Cogn.* 25: 1161–1176.

Ebbinghaus is usually credited with the first empirical demonstration of the forgetting function (although see Galton, 1879). Using his famous savings method, Ebbinghaus recorded the amount of time spent relearning an earlier memory series (usually consisting of nonsense syllables). In one case, he learned eight series of 13 syllables each (to a criterion of two errorless recitations) and then attempted to relearn the same material after one of seven delays ranging from 1 h to 31 days. The percentage savings, calculated as the difference in time spent during initial learning and relearning expressed as a percentage of the original learning time, dropped systematically over delay in a form resembling that shown in **Figure 1**. He expressed some surprise at the form of forgetting, particularly the lessening effect shown at the later delays. Forgetting "in the latter intervals is evidently so slow that it is easy to predict that a complete vanishing of the effect of the first memorization of these series would, if they had been left to themselves, have occurred only after an indefinitely long period of time" (Ebbinghaus, 1885/1964: 76). The question of whether information, once encoded, ever truly vanishes completely (in the absence of relearning or reexposure) remains an issue of concern today (Bahrick, 2000; Wixted, 2004a).

2.12.3.1 Forgetting's Mathematical Form

Psychologists have struggled to characterize the forgetting function in more precise mathematical terms. Ebbinghaus suggested a type of logarithmic function, but many candidates are viable. Linear functions can probably be ruled out, along with simple exponentials, but it is possible to salvage most functions with the right set of assumptions. Psychologists usually choose power – logarithmic, exponential power – or hyperbolic functions, but the decision is often driven by theoretical rather than empirical concerns. At present, despite over a century of effort, no firm consensus on the matter has arisen, although currently many researchers lean toward some kind of power function (e.g., Wixted and Carpenter, 2007).

The failure to reach consensus about forgetting's mathematical form is understandable; the enterprise is fraught with difficulties. For example, one needs criteria for choosing one function, or forgetting model, over another. It is possible to evaluate competing functions simply on the basis of a goodness-of-fit measure, which has typically been the criterion of choice (e.g., Anderson and Schooler, 1991; Wixted and Ebbesen, 1991; Rubin and Wenzel, 1996). However, goodness-of-fit measures usually ignore other important factors, such as the complexity of the function and its psychological viability (Roberts and Pashler, 2000; Pitt et al., 2002; Lee, 2004). There are also serious measurement concerns. The assessment of forgetting requires tracking performance through different points along the measurement scale. Can we really be certain that a drop in retention from, say, 90% correct to 80% correct means the same thing psychologically as a drop from 20% to 10%?

There is also enormous imprecision in the existing data. Forgetting experiments are hard to conduct. Longitudinal studies require testing the same individual at different retention intervals, which is practically difficult and can introduce a testing, or repeated retrieval, confound. In fact, repeated testing of the same information can actually improve overall performance under some conditions (e.g., hypermnesia, Roediger and Challis, 1989). Cross-sectional studies present similar practical difficulties and provide no assurances that the average retention estimates for the different groups accurately represent how forgetting proceeds in an individual (Rubin and Wenzel, 1996; Chechile, 2006). The net result is that most studies report only a handful of retention intervals and, for obvious reasons, the longest sampled retention interval rarely provides any true estimate of memory's permanence. Some memories may remain intact over a lifetime (Bahrick, 1984, 2000) which, in turn, places important constraints on the retention functions that can apply (Wixted, 2004a; Chechile, 2006). Similarly, one cannot ignore performance when the retention interval approaches zero (i.e., immediate testing); some functions are ill-defined at this point, and it is conceivable that special short-term or working memory systems complicate the retention function at very short retention intervals (Chechile, 2006).

There is also the issue of the proper retention measure. As noted, Ebbinghaus measured retention through savings in relearning, but there are many other measurement tools. One can assess memory through proportion correct recall, the d' discriminability index in recognition, or through various indices of priming in implicit or indirect retention measures. The retrieval environment can also be enriched through the introduction of retrieval cues or degraded through the presence of other concurrent tasks. In addition, it is unclear whether delay should be defined as the simple passage of time, time calculated in terms of some kind of relative index (Bjork and Whitten, 1974; Baddeley, 1976), or perhaps the number or quality of intervening events (Waugh and Norman, 1965). Time and events are usually confounded because the passage of time is highly correlated with the number of intervening events (Chechile, 1987).

Despite these difficulties, when an empirical forgetting function is obtained it virtually always resembles the one shown in **Figure 1**. We forget rapidly at first and then retention slowly levels off. Empirically, the proportional rate of forgetting also slows over time, as aptly expressed in Jost's famous law of forgetting: "Given two associations of the same strength, but of different ages, the older falls off less rapidly in a given length of time" (Jost, 1897: 472). Younger memory traces, at least on average, are more vulnerable to the deleterious effects of time than older traces. Note that Jost's law conflicts with simple exponential forgetting functions, which assume constant proportional loss. In an exponential function, the proportional rate of loss remains constant (retention falls by 50% between t and $t+\mathrm{k}$, where k corresponds to the function's half-life), so associations of equal strength, regardless of their age, should decline subsequently at the same rate. They do not, and this places important constraints on the possible mechanisms that underlie forgetting.

2.12.4 Determinants of Forgetting Rates

It is possible to reconcile Jost's law with exponential forgetting if we assume that forgetting rates vary with degree of original learning (Simon, 1966; Wixted, 2004a). For example, if well-learned information is forgotten more slowly than poorly learned information – that is, if the half-life of a memory trace varies with the degree of initial learning – then Jost's law still holds. The possibility that we lose information at a rate determined by its initial strength itself seems imminently reasonable, even intuitive, but it has received little empirical support in the laboratory. Variables that affect acquisition – e.g., word frequency, meaningfulness, similarity, and so forth – typically have little, if any, impact on subsequent forgetting rates (Underwood, 1964; Keppel, 1968).

Slamecka and McElree (1983) allowed the degree of original learning to vary and then assessed retention across several delays. In each case, despite wide differences in original acquisition level, nearly equivalent forgetting slopes were obtained. Importantly, this conclusion held across different retention measures, including free recall, cued recall, category recall, gist recall, and recognition (see **Figure 2**). The same conclusion holds when different mnemonic components or processes are assessed. McBride and Dosher (1999) used Jacoby's process dissociation technique to examine forgetting functions for conscious and automatic components of memory, as a function of the depth of initial processing (semantic vs. graphemic); similar forgetting rates were found for each component, despite differences in the overall level of availability. Even individual difference variables, such as age (Salthouse, 1991) and neurological status (Christensen et al., 1998), commonly fail to produce stark differences in either the form or the rate of forgetting.

These data, among many others, suggest that acquisition and forgetting are not merely two sides of the same coin; rather, variables that affect acquisition may not affect the forgetting process. Forgetting proceeds in its characteristic way, regardless of the retention measure or initial acquisition level, much like the action potential of a neuron, once generated, travels forward in a characteristic (all-or-none) fashion. Of course, this conclusion must be hedged a bit for all the methodological considerations that have been listed; plus, controversies have raged over how to measure the loss of information over time properly (e.g., Slamecka, 1985; Bogartz, 1990; Loftus and Bamber, 1990). Still, it is generally conceded that forgetting rates remain invariant as acquisition variables, mnemonic processes, and measurement vehicles are manipulated, although exceptions can be found in the literature. For example, McDonald and colleagues (2006) recently found evidence for acquisition-based heterogeneity in forgetting when a novel statistical procedure, multilevel modeling, was applied to the retention data and certain other methodological concerns, such as the mnemonic strategy adopted by the participant, were controlled. The debate continues, as it has for the past century.

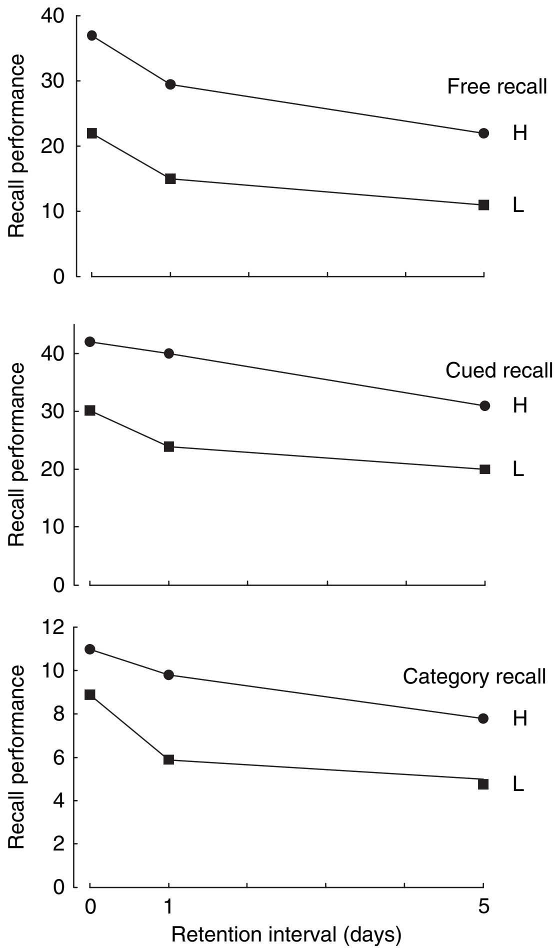

Figure 2 Forgetting functions obtained for high (H) and low (L) degrees of original learning as revealed for three different retention measures. After Slamecka NJ and McElree B (1983) Normal forgetting of verbal lists as a function of their degree of learning. *J. Exp. Psychol. Learn. Mem. Cogn.* 9: 384–397.

2.12.5 Mechanisms of Forgetting

It may be difficult to characterize the forgetting function completely, but every researcher recognizes that memory changes systematically with time. As noted earlier, there are excellent reasons to believe that forgetting is adaptive and, more importantly, that the forgetting function mimics the way events occur and recur in the natural world. Consequently, it is reasonable to search for a mechanism, or set of mechanisms, that affects the availability of learned material. Of course, any acceptable theoretical account of why we forget must come to grips with the regularity of the forgetting function itself. One cannot simply argue, for example, that forgetting is cue-dependent – that is, we forget in the absence of an appropriate retrieval cue – without also explaining why the form of forgetting is so regular and predictable.

Historically, researchers have appealed to three primary causal mechanisms to explain forgetting: autonomous decay, interference from other acquired information, and altered stimulus conditions (McGeoch, 1932; Bower and Forgas, 2000). There is a fourth mechanism, active inhibition induced by retrieval, that has been proposed more recently (Bjork and Bjork, 1992; Anderson, 2003), although it, too, has some interesting historical antecedents (e.g., Freud's concept of repression). Each of these mechanisms is discussed in more detail in the sections that follow.

2.12.5.1 Decay

When memory theorists use the term *decay*, they mean forgetting that occurs spontaneously with the passage of time. Decay is assumed to be autonomous, which means that its progression does not depend on some other active mnemonic process (such as the acquisition of new information). Of course, everyone believes there must be some kind of neurological underpinning for the decay process, such as a metabolic process that erodes or overwrites synaptic connections, but the mechanism itself is left unspecified. The natural process of radioactive decay is sometimes used as a rough analogy, in which a constant proportion of radioactivity is lost in a fixed unit of time.

To make a principled empirical argument for decay, it is necessary to show that forgetting proceeds in the absence of other activities, such as rehearsal or interference. As we discuss in the next section, newly established memories are quite susceptible to interference from other learned material, so it is necessary to control, if possible, for this factor. Rehearsal also needs to be ruled out, because repetition presumably counteracts the deteriorating effects of the decay process. One needs to create, in essence, a kind of mental vacuum in which time, and little else, is allowed to vary. If information is still forgotten in its characteristic way, then a time-based process, decay, must be responsible.

Various strategies have been employed to detect decay. Most have relied on a dual-task methodology in which the participant is asked to perform an attention-demanding task while simultaneously retaining critical target memories. For example, Reitman (1974) instructed participants to retain groups of words while detecting the occurrence of a tone signal in background noise. The tone detection task was assumed to be sufficiently dissimilar to the words to prevent interference, which typically depends on similarity, and sufficiently taxing to prevent rehearsal (although see Roediger et al., 1977). Similarly, Cowan and colleagues (1997) had people compare the frequencies of two tones separated by delays while performing a silent visual tracking task. It is difficult, perhaps impossible, to rehearse a pure tone, and the intervening visual distractor task seemed unlikely to produce interference, so any decline in the tone task was attributed to decay. In both of these cases, retention performance did decline with delay, implicating decay (although see Cowan et al., 2001).

Note that each of these paradigms investigated retention over the short term, that is, over a time-course of seconds. In fact, the concept of decay is rarely used to explain long-term forgetting, due in large part to the seminal arguments of John McGeoch (1932). Among other things, McGeoch noted that memories often remain highly available or even improve over time, a fact that seems fundamentally incompatible with decay. For example, when people are given repeated opportunities to recall earlier-presented material, they will often recall items on the second or third attempt that they failed to recall initially, a phenomenon known as reminiscence (Ballard, 1913). The Pavlovian phenomenon of spontaneous recovery, in which an extinguished conditioned response recovers after a period of rest, is another case in point (for empirical evidence of spontaneous recovery in human participants, see Wheeler, 1995).

McGeoch (1932) further noted that when the passage of time is held constant, the amount of

forgetting depends on the specific activities that occur during the delay. After learning a list of verbal items, if one group of participants rests and a second group learns another list, recovery of the initial list will be sharply impaired in the second group (Müller and Pilzecker, 1900). Moreover, the amount of forgetting will depend on similarity between the original material and the interpolated material (see Osgood, 1953). Forgetting is correlated with the passage of time, McGeoch argued, but the relationship is not causative. It is the events that happen 'in time' that produce forgetting, not time itself.

In the case of short-term retention, however, the situation is somewhat different. Decay remains quite popular among theorists in this arena because short-term retention is thought to tap special storage systems. In the working memory model of Baddeley and Hitch (1974), for example, information is retained for brief intervals in the service of more complex forms of processing (e.g., language), and special 'activity traces' are assumed to be stored in various loops, buffers, and sketchpads (Baddeley, 2000). In the absence of rehearsal, which serves a refreshing function, these traces decay autonomously. Most working memory theorists accept that short-term forgetting can occur as a consequence of other means, such as interference, but decay is assigned an important and even pivotal role (e.g., Page and Norris, 1998).

Memory traces may indeed decay in some circumstances and not in others. There could be something special about short-term activity traces, those engendered and maintained by special short-term memory systems, but the concept remains controversial. For example, Nairne (2002b) has shown that each of McGeoch's main arguments against decay in long-term memory apply equally well to immediate retention: Short-term retention can decline, remain the same, or even improve over time depending on the circumstance. To illustrate, in a study by Turvey et al. (1970), different groups of people were asked to count backward as a distractor activity for group-specific intervals prior to recall (e.g., one group counted for 10 s, another for 15 s, and another for 20 s). Equivalent amounts of forgetting were found across groups in this between-subject design (0.33, 0.30, and 0.30, respectively) (see also Greene, 1996). Moreover, on a critical trial all groups were switched to the same 15-s distractor period. Retention performance dropped in the 10-s group (from 0.33 to 0.20), stayed roughly constant in the 15-s group (0.30 to 0.28), and improved in the 20-s group (0.30 to 0.38). Note that the passage of time – and therefore the opportunity for decay – was equated across the groups on the critical 15-s trial, yet performance depended significantly on the timing of prior trials.

Space does not permit a complete review of all the relevant studies. Suffice to say that the correlation between time and forgetting is far from perfect even when retention is tested over intervals lasting seconds. Importantly, however, such findings by themselves do not rule out the concept of decay. It is still possible that memory traces decay with time, but particularly supportive retrieval environments can counteract the loss. In the Turvey et al. (1970) study, for instance, moving from a 20-s distractor-filled retention interval to a 15-s interval may have helped the participant discriminate to-be-remembered items from information recalled on prior trials (Nairne, 2002b). Similarly, the emergence of newly recalled items on a second or third recall attempt could simply reflect participants' ability to use recalled items as cues to help them recall new items. The state of the trace could still be degraded at time 2 relative to time 1, but a more supportive retrieval environment at time 2 nets improved recall.

At the same time, recognizing that retention cannot be predicted from the state of the memory trace, without considering the retrieval environment, seriously undermines the theoretical utility of decay. Appealing to decay as the source of forgetting is like appealing to strength as the source of remembering. As Endel Tulving (1983) has argued, memory traces "do not have strength independently of the conditions under which they are actualized" (Tulving, 1983: 240–241). Thus, losing trace features over time may or may not impair retention; it will depend on the particular features that are lost and their compatibility with the retrieval cues present at the time of retention testing (Nairne, 2002a).

2.12.5.2 Interference: Trace Degradation

The second, and more popular, interpretive tool used to explain forgetting is *interference*. The concept of interference is multifaceted, having several distinct meanings, but the important common denominator is the occurrence of other mnemonic events. We forget because other events interfere with the storage or recovery of target memories. Unlike decay, the interference perspective assumes that if one could create a mental vacuum – that is, if you could measure the state of a memory trace over time in the absence of other mnemonic events or activities – there would be no decline in the integrity of the trace. Forgetting

occurs because other events or activities, particularly ones that are memory-based, happen 'in time.'

There are two ways that interference is thought to operate. First, newly learned material can overwrite, erase, displace, or otherwise degrade an existing memory trace. The details are usually left unspecified, although it is generally assumed that similarity between original and new learning increases the extent of the interfering effect. As noted above, many studies have shown that if the retention interval is held constant, the nature of the activities that occur between study and test importantly determines what and how much is forgotten. In a classic study by Jenkins and Dallenbach (1924), for instance, people recalled more information if they slept through a retention interval than if they remained awake; the assumption here is that sleep protects one from the potentially damaging effects of interpolated interference. Comparable findings occur even for amnesic subjects: After hearing words or stories, if amnesic patients are allowed to spend a retention interval in a dark and quiet room, their subsequent retention is vastly improved relative to an interference control (Cowan et al., 2004).

Moreover, in an interesting parallel to decay, some researchers assume that the damaging effects of interference depend on the passage of time as well. Rather than exerting a negative effect, however, memory traces are assumed to become less vulnerable to the effects of interference with time because of a trace consolidation process. The notion that memory traces consolidate is widely accepted by neuroscientists, partly because retrograde amnesia, the loss of memories formed prior to brain damage, shows a distinct temporal gradient (Ribot, 1881). A blow to the head is more likely to lead to the loss of recently formed memories than to the loss of memories from the more distant past. Presumably this pattern occurs because the older traces have consolidated and, as a consequence, are less susceptible to interference. As Wixted (2004a,b) has recently observed, exactly the same reasoning can be applied to general mnemonic principles such as Jost's law: Given two traces of the same strength, but of different ages, retention of the older one will fall off less rapidly in a given length of time, presumably because the older trace has sufficiently consolidated.

In the laboratory, however, the concept of consolidation has a more checkered past. In fact, it has been largely rejected by memory theorists for decades because laboratory-based experiments typically fail to show convincing temporal gradients. It is possible to obtain robust retroactive interference – the term used to describe interference arising from events occurring after the target memory is established – when the interfering event occurs days or even weeks after the original encoding, a period far exceeding any reasonable consolidation time. Moreover, the interference that is obtained after a short delay, when consolidation processes are presumably active, is usually comparable in size to the interference obtained after lengthy delays (see Wickelgren, 1977, for a review).

In fairness to consolidation theory, though, the fact that robust retroactive interference occurs after a long delay does not rule out a consolidation process, for much the same reason that retention after a lengthy delay does not rule out a decay process. As discussed in the next section, interpolated learning can easily decrease the accessibility of a fully formed memory trace by impairing the diagnostic value of an associated retrieval cue. Moreover, Wixted (2004b) has recently challenged the accepted dogma concerning temporal gradients, arguing that traditional retroactive interference designs introduce methodological problems that cloud interpretation. Also, some clinical cases indicate that consolidation may last a relatively long time (Dudai, 2004).

Still, it is important to recognize that consolidation, even if empirically verified, can never stand as a completely adequate account of forgetting (see also Meeter and Murre, 2004). Consolidation theory, like decay theory, essentially reduces to a set of claims about how the integrity of a memory trace changes with the passage of time. As noted, memory traces do not have retention strength outside of the conditions under which they are accessed. For any given trace, degraded or otherwise, there are presumably retrieval conditions that will promote or hinder successful retrieval. To explain forgetting or retention adequately, one needs to consider the state of the memory trace as well as the conditions present at the time retrieval is attempted (Tulving, 1983).

2.12.5.3 Interference: Cue Impairment

We mentioned earlier that interference is thought to operate in two ways. The first, just discussed, is the degrading effect that new learning can have on the integrity of an existing trace. The second route places the locus of interference not in the trace itself, but in its eliciting retrieval cue. Psychologists generally assume that remembering is cue-driven. Memories are not thought to arise spontaneously; instead, they are activated by the presence of associated retrieval

cues. For decades, the empirical paradigm of choice among researchers was paired-associate learning in which participants are asked to associate target words (e.g., 'town') with cue words (e.g., 'cart'). The advantage of paired-associate learning is that it allows the experimenter to test the integrity of a specific memory by cuing the participant with its associated retrieval cue. The ability to reproduce the target word in the presence of its linked cue is used as the index of retention. When people fail to produce the appropriate response, given the presence of the retrieval cue, then obviously forgetting has occurred.

Once again, forgetting in such a context could occur because the integrity of the target memory has been degraded, either through a decay process or as a by-product of subsequent activity. However, it is also possible that the cue–target association is impaired, leaving the integrity of the target trace intact. In such a case, forgetting occurs because the retrieval cue is unable to elicit or reproduce a previously associated target memory. There are two reasons why the effectiveness of a retrieval cue can become impaired. First, as suggested by early interference theorists (e.g., Melton and Irwin, 1940; McGeoch and Irion, 1952), subsequent activity might lead to 'unlearning' of the cue–target association. Suppose, for example, that after learning an association between 'cart' and 'town', the cue 'cart' is subsequently associated with other targets (e.g., 'block' or 'train'). During this relearning phase, 'cart' occurs in the absence of 'town' and, therefore, the 'cart–town' association extinguishes, much like the process of extinction in Pavlovian conditioning (Pavlov, 1927).

The second and more commonly accepted mechanism for cue impairment is target competition. Even in the absence of unlearning, if a retrieval cue becomes linked to several targets, its ability to elicit any one of those targets lessens. So, if 'cart' is associated initially with 'town,' but then is later paired with 'block' or 'train,' the probability that 'town' will be produced with 'cart' on demand declines (see **Figure 3**). This characteristic of retrieval cues is known by several names, including *response competition, cue distinctiveness,* the *fan effect* (where 'fan' refers to the number of associated target responses), and *cue overload.* Historically the term 'response competition' was used in conjunction with 'unlearning' to form the two 'factors' of the famous two-factor theory of forgetting (McGeoch, 1942; Postman, 1961). Currently, the more popular moniker is *cue overload*, for reasons that will be discussed momentarily.

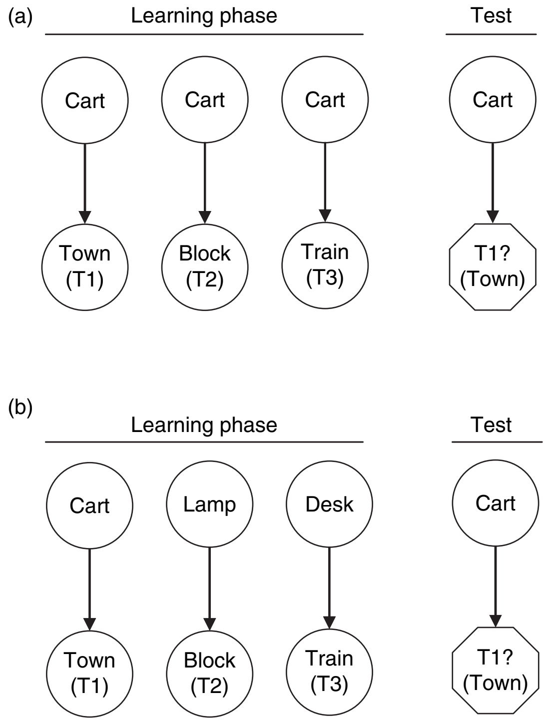

Figure 3 Illustration of a cue overload situation. In (a) the retrieval cue is linked with several targets, whereas in (b) it is associated only with one target. The predictive value of the cue in situation (a) is lower than in (b) due to cue overload.

The basic phenomenon of cue overload is well supported empirically. For example, as the number of study items from a particular category increases, the category name becomes a less effective cue for eliciting any one item in particular (Tulving and Pearlstone, 1966; Roediger, 1973). Cue overload explains the list length effect as well: recall of any given item from a memory list declines as the length of the list increases. The list length effect occurs, presumably, because people use some representation of the list itself as a cue and it becomes 'overloaded' (i.e., less diagnostic of any particular item) as the number of items subsumed under the list cue increases (Watkins and Watkins, 1975). More directly, it is well established that increasing the number of associated responses to a given target – that is, increasing the cue's associative 'fan' – slows down people's ability to verify any particular cue–target pairing (Anderson, 1974; Anderson and Reder, 1999).

Cue overload is the preferred term partly because the locus of interference is believed to lie primarily in the target selection phase rather than in competition among already selected responses. Although this distinction may seem somewhat arbitrary, or at least model-specific, its roots lie in classic work by

Barnes and Underwood (1959) on the modified modified free recall test (MMFR). In an MMFR test, people are asked to recall any or all response terms that have been associated with a cue. There is no requirement to recall a specific response term which, in turn, presumably eliminates any response competition (because any and all responses can be produced). Significant interference is still obtained in the MMFR test – that is, one's ability to recall 'town' to 'cart' remains impaired if other responses have been associated with the cue – and this fact was used by interference theorists to support the concept of unlearning (Postman, 1961). Modern theorists rarely invoke unlearning *per se*, choosing to argue instead that new associations essentially block or impair access to old ones during the target selection phase, or that the recall of one item leads to inhibition or suppression of the other (see Anderson and Neely, 1996, for a review).

Ascribing interference effects to cue impairment – specifically, the ability of a cue to produce a specific target – has considerable advantages. For example, it enables the theorist to explain both retroactive and proactive interference with a single mechanism. As discussed earlier, retroactive interference occurs when newly learned material acts retroactively to impair earlier learning; proactive interference occurs when information learned at time 1 interferes with the ability to access information learned at time 2 (see **Figure 4**). Underwood (1957) provided convincing evidence that much of what is forgotten in standard verbal learning experiments can actually be attributed to proactive interference (i.e., prior learning). **Figure 5**, which is based on data compiled originally by Underwood, shows proportion correct recall of a serial list, tested after an unfilled 24-h interval, as a function of the number of lists learned previously in the experimental context. When the critical list is the only list learned, about 80% of the material will be retained after 24 h; as the potential for proactive interference increases – that is, as the number of prior lists learned increases – delayed retention drops precipitously.

The robust nature of proactive interference is troubling for most of the forgetting mechanisms we have discussed. Why should learning 'cart–town' at time 1 impair one's ability to recover 'cart–block' learned at time 2? Certainly neither decay, nor overwriting, nor unlearning can explain the phenomenon because each is ascribed to things that happen after the point of acquisition. Significant proactive interference is found on an MMFR test as well (Koppenaal, 1963),

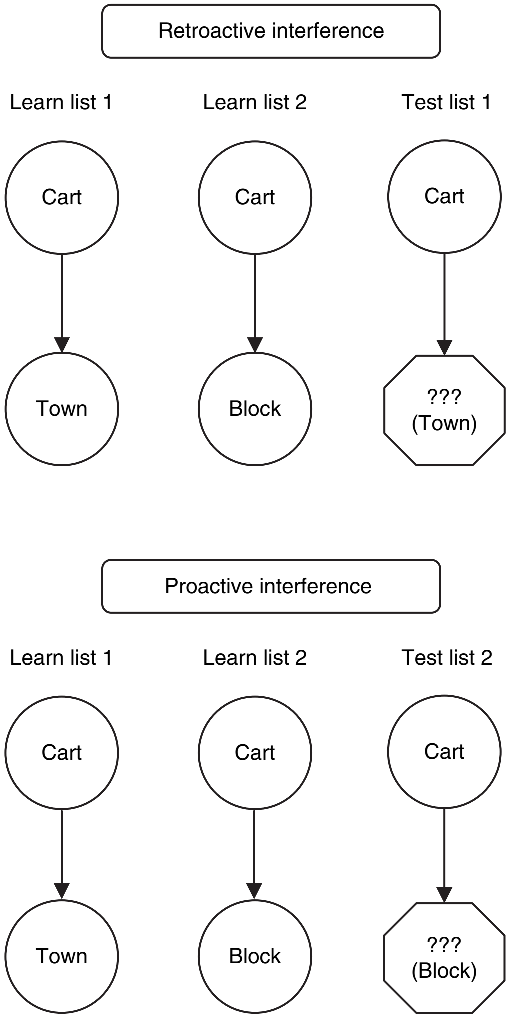

Figure 4 Representation of experimental settings revealing retroactive interference and proactive interference.

which rules out historical versions of response competition. The only viable forgetting mechanism is cue overload: As the number of targets associated with a cue increases, the ability of the cue to access any particular target declines. Note there is nothing about the order of acquisition that is inherent in the concept of cue overload, although the availability of particular cue–target associations does change in complex ways with time and testing method (see Postman and Underwood, 1973).

2.12.5.4 Cue Availability

In addition to interference from events that happen in time, forgetting can also arise from altered stimulus conditions, that is, when "the stimuli necessary to

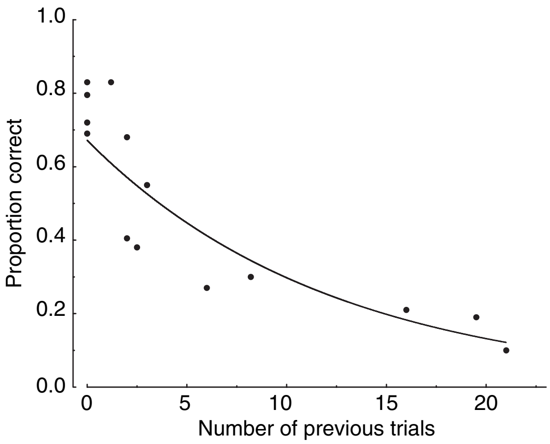

Figure 5 Proportion of correct responses plotted as a function of the number of previous trials. Data were collected from 14 different studies. After Underwood BJ (1957) Interference and forgetting. *Psychol. Rev.* 64: 49–60.

elicit the originally learned acts are not effectively present" (McGeoch, 1942: 501). Endel Tulving (1974) coined the term *cue-dependent forgetting* to describe the situation. Put simply, if you have learned to associate the target 'town' with the retrieval cue 'cart,' then in the absence of 'cart' you are unlikely to remember 'town.' In this case, there is no degradation or impairment of the target trace, nor any impairment in the cue-target association: We forget simply because we lack the right retrieval cue.

Cue-dependent forgetting helps to explain why memories that seem to have been lost can reappear at a later time. To use the terminology of Tulving and Pearlstone (1966), information can be 'available' in memory, in the sense that the encoded information remains intact somewhere in the memory system, but not 'accessible.' Accessibility requires an appropriate retrieval cue which may be absent at time 1 but may reappear at time 2. Of course, the fact that forgetting proceeds in such a regular fashion, with such a characteristic retention function, suggests that the appearance of retrieval cues may be time-locked. Indeed, psychologists have used the concept of context, which is assumed to change systematically over time, to explain the regularity of forgetting. During presentation, information becomes associated with the context which then changes in accordance with the flow of normal activities (see Estes, 1955; Mensink and Raaijmakers, 1988; Brown et al., 2000, for some specific theories on how context changes). (Note that the examples discussed so far involve external cues (e.g., presenting a cue previously associated with a target), but cue-dependent forgetting can also be demonstrated with internal cues (e.g., mood state; Bower, 1981). In this case, an internal stimulus (such as a mood state) acts as the functional retrieval cue.)

Accepting that remembering (and forgetting) is cue-dependent encourages us to specify the conditions that determine (1) when relevant retrieval cues will be available and (2) the effectiveness of those cues when present. Psychologists have spent decades investigating the second point and have reached consensus that retrieval cues are effective to the extent that they 'match' the contents of the original encoding. The encoding specificity principle, first articulated by Tulving and Thomson (1973), states that retrieval cues will be effective in eliciting targets if and only if the information about them and their relation to the to-be-remembered target is stored at the time of encoding. Thus, the conditions of encoding will uniquely determine whether any given cue will be effective in recovering a prior target episode. Retrieval cues will work if and only if they match, and consequently are a part of, the original encoding complex.

The encoding specificity principle asserts that preexisting associations between cues and targets, such as the semantic relationship between 'bloom' and 'flower,' cannot be used *a priori* to predict the effectiveness of one cue for another. This runs counter to intuition because one would normally expect a strong associate to elicit the target naturally, allowing it to be confirmed easily as a member of the study list. Yet, Tulving and Thomson (1973) showed that a weak associate to the target 'flower,' such as 'fruit,' can actually be a better cue than the strong associate when conditions promote the encoding of the weak associate during study. It follows as well that retention will depend on the match between the retrieval cue and the target memory, as encoded, rather than on the cue and target as originally presented. Presenting the same nominal cue at test will not necessarily be effective. It will depend on whether the participant interprets the retrieval cue at test in the same way that he or she interpreted the cue during the original encoding.

Virtually all psychologists recognize the importance of matching the encoding and retrieval environments, to assure the availability of an appropriate cue, but some have questioned the role of the match *per se*. For example, Nairne (2001, 2002a) has argued that it is the diagnostic value of the retrieval cue that really matters. Rather than a passive matching process, retrieval is better characterized as an active selection process wherein cues are used to

pick and choose from among viable retrieval candidates. Matching the retrieval cue with the original encoding context, as encoded, can be expected to increase the diagnostic value of the cue in most situations, but it is easy to conceive of situations in which increasing the match will not improve retention, or perhaps even lower it. For example, if features are added to the cue that match the target exactly, but also match additional non-target items, then the difficulty of the target selection process can increase and performance decline.

The situation is somewhat akin to the relationship between stimulus intensity and the perception of brightness. It is generally the case that increasing the intensity of a light source makes things look brighter, but what mainly determines brightness perception is relative intensity information. It is the number of photons falling in a given spot relative to the number falling in surrounding spots that determines how bright the central spot appears. In fact, it is possible to increase the absolute light intensity falling on the spot and make it look darker (as long as light intensity in the surround is greater still). In discussing brightness perception, it is misleading to focus on light intensity *per se* because our visual system tends to throw away absolute information in favor of relative comparisons. Similarly, for retrieval, it is misleading to focus on the absolute match between a retrieval cue and an encoded target when it is actually the diagnostic value of the cue – the relative match – that really matters (see Nairne, 2002a).

2.12.5.5 Retrieval-Induced Inhibition

One interesting feature of the forgetting mechanisms that we have discussed so far is their passive nature. Things happen to the memory trace which yields it less recoverable: It decays, it is degraded by subsequent events, the cue–target association is extinguished by new learning, or the memory target exists in an unrecoverable state because an appropriate retrieval cue is lacking. Yet from an adaptive perspective, it seems likely that our memory systems may have evolved active mechanisms to inhibit or suppress information in specific situations in which that information is not needed. Inhibitory processes certainly play an important role in the nervous system, particularly in neural communication, so it is not a stretch to assume that inhibition is vital to memory processing as well. We might also assume that the effects of inhibition, when it occurs, must be temporary given the continuously changing nature of our processing goals (MacLeod and Macrae, 2001).

To be clear, in modern memory theory the concept of inhibition is synonymous with suppression. When a memory trace is inhibited, it is not degraded, damaged, or impaired, it is simply rendered temporarily unavailable by an active suppression process. Importantly, suppression of this sort differs from the retrieval blocking produced as a consequence of cue overload. In cue overload, competition among viable targets produces a 'winner' and unselected targets suffer as a consequence, but there is no need to assume suppression of the 'losers.' Likewise, recall is often claimed to have inhibitory properties (Roediger, 1974, 1978) because the act of recalling one item can lower the probability that other items will also be recalled; however, this kind of 'output interference' is generally assumed to result from biased sampling rather than from suppression. Once an item is recalled, the probability that it will be sampled again increases (it is primed) which, in turn, lowers the probability that other targets will be sampled. Note that inhibition might well occur as a by-product of cue overload, or biased sampling in recall, but it is not needed to produce forgetting in these instances.

As it turns out, the best evidence for inhibition comes from an empirical procedure known as the *retrieval practice paradigm* (Anderson et al., 1994). Here, people first learn lists of category-exemplar pairs (e.g., 'fruit–banana,' 'drink–scotch') and are then asked to practice retrieval of half of the exemplars from half of the list categories. Practice takes the form of completing stem-recall tests ('fruit-or____?') which people are required to complete several times. Last, after a short delay, a final category cued recall test is given for all of the exemplars (see **Figure 6**). There are two main findings of note: First, recall of the practiced exemplars is superior to that of the unpracticed exemplars; second, recall of the unpracticed exemplars from the practiced categories is impaired relative to exemplars from the unpracticed categories. Thus, practicing the recall of 'fruit–orange' impairs recall of the unpracticed exemplar 'fruit–banana,' below the baseline recall level for exemplars from the unpracticed categories. This impairment is known as *retrieval-induced forgetting* and is thought to accrue from an active inhibitory process.

Of course, there are other interpretations of these findings. For instance, one could appeal simply to retrieval blocking. Practicing 'fruit–orange' increases the strength of the cue–target association which, in turn, should bias the system to sample 'orange' in the

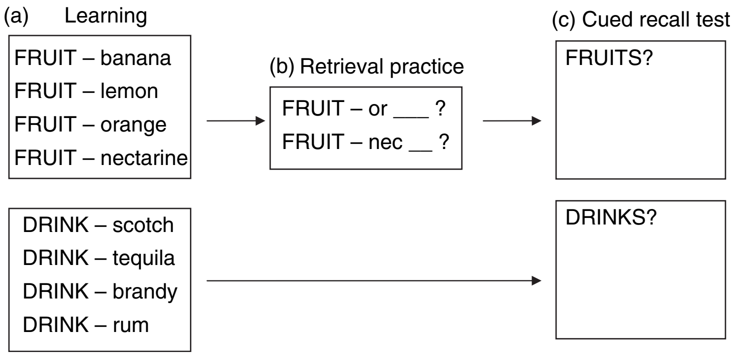

Figure 6 A simplified version of the retrieval practice paradigm. (a) Elements of two different categories are learned. (b) Half of the elements from one category are subject to a retrieval practice phase–cued word stem completion. (c) After a distractor period, participants respond to a cued recall test with the category name; participants are required to recall all the information learned in (a).

presence of 'fruit' to the exclusion of other exemplars. What supports the presence of active inhibition is the finding that the impairment is cue independent, that is, recall of 'orange' is impaired even when queried with cues unrelated to the category cue 'fruit' (see Anderson, 2003). In retrieval blocking, the impairment results from competition among targets elicited or matched by a given cue which, in turn, should lessen if 'orange' is cued by a target that is less overloaded. As a result, the retrieval practice paradigm seems to create a target that is truly suppressed, leading to impairment regardless of how the target is queried at the point of test.

As noted earlier, the idea that our memory systems may have developed mechanisms for actively suppressing information makes considerable adaptive sense. For any given constellation of retrieval cues, there is likely to be a wide array of potentially recallable responses, so it is in our interest to reduce irrelevant clutter. A telephone, for example, potentially elicits dozens of numbers to call, but we focus on the number at hand and push the remaining numbers out of mind (Levy and Anderson, 2002). A similar task faces us in perception: We must focus our attention on relevant portions of the sensory/perceptual message and block out the irrelevant ones. In memory, as in perception, it is adaptive to exercise cognitive control in our efforts to prioritize functioning. Inhibition – i.e., active suppression – is undoubtedly a useful weapon in the arsenal of cognitive control.

2.12.5.6 Motivated Forgetting

Proposals about inhibitory control in memory retrieval lead one naturally to the concept of repression, Freud's proposed defense mechanism (Freud, 1915). Although exactly what Freud meant by repression is open to some interpretation (Boag, 2006), it is generally conceived as a mechanism for preventing anxiety-inducing memories, usually traumatic, from entering the sphere of conscious awareness (Gleaves et al., 2004). Motivated forgetting of this sort seems adaptive: Unpleasant or traumatic memories can interfere with normal functioning, and the associated stress reactions can lead to long-term health consequences as well.

Adaptive speculations aside, is there solid empirical evidence for repression? Unfortunately, most of the relevant data are anecdotal. For obvious reasons, laboratory-based investigations of trauma-induced memory suppression are virtually impossible to conduct. The case for repression continues to rest largely on the many reports of trauma-based amnesia, and subsequent recovery during therapy, that have been obtained in clinical settings. Some relevant survey data and/or clinical cases exist as well, during which people with known histories of sexual abuse have reported periods of amnesia for their abuse (e.g., Williams, 1994; Schooler et al., 1997). However, perhaps not surprisingly, few in the scientific community find these data to be particularly convincing (e.g., Kihlstrom, 2004).

The clinical data are controversial for several reasons. First, in therapeutic settings it is often difficult to verify whether the traumatic event actually occurred or occurred in the form revealed by the recovered memory. Therapists usually do not seek independent corroboration of their clients' reports, again largely for ethical reasons (Shobe and Kihlstrom, 2002). Second, many psychologists believe that therapist–client interactions are particularly prone to the

inducement of false memories, perhaps because prior abuse is believed by many to be an important determinant of psychological problems. This is not to imply that false recovered memories are purposely implanted by the therapist, but they can occur as an unintentional by-product (e.g., Porter et al., 1999).

Finally, even if the traumatic event did occur, and the recovered memory is accurate in all details, this does not mean that an active repression process has produced the forgetting. As documented in this chapter, there are many reasons why people forget, and normal forgetting processes could easily account for many, if not all, of the verifiable cases of repressed memory. Just because a memory is traumatic does not mean that it is insensitive to decay, interference, or cue-dependent forgetting. In fact, given that many clinically relevant instances of abuse are thought to occur during childhood, or are accompanied by considerable emotional distress, it is perhaps not surprising that relevant retrieval cues are sometimes lacking at later points in time. Inhibitory mnemonic processes probably do exist, as evidence from the retrieval practice paradigm has shown, but whether there are special inhibitory processes (i.e., repression) that apply when memories are traumatic or emotionally tinged remains speculative at best.

2.12.6 Conclusions

Forgetting occurs when we fail to recover information that has been experienced previously. As noted initially, the common tendency is to label forgetting as a nuisance (or worse), but the process itself is actually quite adaptive. Imagine entering the parking garage after work and simultaneously recovering the locations of all of your previous parking spots. In his famous case study of the Russian journalist S., who was plagued by an inability to forget, Luria (1968) describes the torment S. experienced daily. S. had great trouble reading books, for example, because words and phrases so flooded his mind with previous associations that he was unable to concentrate. To avoid a truly cluttered mind, it is reasonable to assume that forgetting is a design feature of memory, that is, a cognitive capability that was selected for and maintained during the evolutionary history of our species.

Given the role that forgetting plays in normal functioning, it is reasonable to assume as well that there are many routes to forgetting. In the bulk of this chapter, we discussed a variety of forgetting mechanisms, everything from decay to interference to inhibition (summarized in **Table 1**). Each mechanism carries some weight of evidence and continues

Table 1 Summary of the mechanisms of forgetting presented in the chapter

Mechanism	*Source*	*Process*	*Effect*
Decay	Time	Autonomous process	Traces, or trace features, are permanently lost with time
Consolidation	Time	Autonomous process	Memories become more resistant to forgetting with time
Interference	Other acquired information	Trace degradation (Retroactive interference)	Newly acquired information damages the integrity of existing memory traces (e.g., overwriting)
		Cue impairment (Retroactive interference) (Proactive interference)	Other acquired information impairs the ability of a cue to produce a specific target due to the unlearning of preexisting associations between the cue and the target, or to target competition resulting from the association of the same cue with several targets (cue overload)
Cue availability	Altered stimulus conditions	Absence of an effective retrieval cue	The cue needed to elicit the originally learned information is not available
Retrieval-induced inhibition	Retrieval of other information	Active inhibition process	Temporary inhibition of information in situations in which it is not needed or when it competes with other target memories
Motivated forgetting	Repression	Active repression	Prevents anxiety-inducing memories, usually traumatic, from entering the sphere of conscious awareness

to have many advocates. It would be improper to conclude that any one of these mechanisms is 'the' forgetting mechanism because different situations will undoubtedly demand different forgetting solutions. In some circumstances, important information needs to be suppressed temporarily; in others it may be in our interest to forget things permanently (or nearly so).

Whatever the mechanism, however, it does remain a challenge for memory theorists to explain the regularity of the forgetting function. As documented earlier, the function relating recovery to time is very regular in form (see **Figure 1**). Attributing forgetting to interference from subsequent events, or to the action of cue overload, does little to explain why the forgetting curve is consistently negatively accelerated. Forgetting may indeed be cue-dependent, but then what determines the availability of appropriate cues? As the forgetting curve informs, forgetting is by no means a random occurrence.

Lastly, it is worth noting that forgetting can never be measured directly: We can only measure what has been 'remembered,' at a particular time, given a particular context. And, as we have described, what appears to have been forgotten may, in fact, turn out to be recoverable at another time or in a different context. To the extent that our memory systems are designed to use the past, in combination with the present, to generate adaptive behavior, then the variability of memory is hardly surprising. Stored information should be retrieved when it is necessary and not otherwise. In this sense it is a mistake to speculate about the ultimate 'cause' of forgetting, or to consider forgetting as a breakdown in normal functioning, because our memory systems are not designed to recover stored information at will: Recovery will depend on the situation and, more importantly, on the particular adaptive problem at hand.

Acknowledgment

J. N. S. Pandeirada was supported by fellowship SFRH/BPD/21534/2005 from the Portuguese Foundation for Science and Technology.

References

Altmann EM and Gray WD (2002) Forgetting to remember: The functional relationship of decay and interference. *Psychol. Sci.* 13: 27–33.

Anderson JR (1974) Retrieval of propositional information from long-term memory. *Cogn. Psychol.* 6: 451–474.

Anderson JR and Reder LM (1999) The fan effect: New results and new theories. *J. Exp. Psychol. Gen.* 128: 186–197.

Anderson JR and Schooler LJ (1991) Reflections of the environment in memory. *Psychol. Sci.* 2: 396–408.

Anderson MC (2003) Rethinking interference theory: Executive control and the mechanisms of forgetting. *J. Mem. Lang.* 49: 415–445.

Anderson MC and Neely JH (1996) Interference and inhibition in memory retrieval. In: Bjork EL and Bjork RA (eds.) *Memory: Handbook of Perception and Cognition*, pp. 237–313. New York: Academic Press.

Anderson MC, Bjork RA, and Bjork EL (1994) Remembering can cause forgetting: Retrieval dynamics in long-term memory. *J. Exp. Psychol. Learn. Mem. Cogn.* 20: 1063–1087.

Baddeley AD (1976) *The Psychology of Memory*. New York: Basic Books.

Baddeley AD (2000) The episodic buffer: A new component of working memory? *Trends Cogn. Sci.* 4: 417–423.

Baddeley AD and Hitch GJ (1974) Working memory. In: Bower GA (ed.) *Recent Advances in Learning and Motivation*, vol. 8, pp. 47–89. New York: Academic Press.

Bahrick HP (1984) Semantic memory content in permastore: Fifty years of memory for Spanish learned in school. *J. Exp. Psychol. Gen.* 113: 1–29.

Bahrick HP (2000) Long-term maintenance of knowledge. In: Tulving E and Craik FIM (eds.) *The Oxford Handbook of Memory*, pp. 347–362. New York NY: Oxford University Press.

Ballard PB (1913) Obliviscence and reminiscence. *Br. J. Psychol. Monogr. Suppl.* 1: 1–82.

Barnes JM and Underwood BJ (1959) "Fate" of first-list associations in transfer theory. *J. Exp. Psychol.* 58: 97–105.

Bartlett FC (1932) *Remembering: A Study in Experimental and Social Psychology*. New York: Cambridge University Press.

Bjork EL and Bjork RA (1988) On the adaptive aspects of retrieval failure in autobiographical memory. In: Gruneberg MM, Morris PE, and Sykes RN (eds.) *Practical Aspects of Memory: Current Research and Issues,* vol. 1, pp. 283–288. London: Wiley.

Bjork EL and Bjork RA (1996) Continuing influences of to-be-forgotten information. *Conscious. Cogn.* 5: 176–196.

Bjork RA and Bjork EL (1992) A new theory of disuse and an old theory of stimulus fluctuation. In: Healy AF, Kosslyn SM, and Shiffrin RM (eds.) *From Learning Processes to Cognitive Processes: Essays in Honor of William K. Estes,* vol. 2, pp. 35–67. Hillsdale, NJ: Erlbaum.

Bjork RA and Whitten WB (1974) Recency-sensitive retrieval processes in long-term free recall. *Cogn. Psychol.* 6: 173–189.

Boag S (2006) Freudian repression, the common view, and pathological science. *Rev. Gen. Psychol.* 10: 74–86.

Bogartz RS (1990) Learning-forgetting rate independence defined by forgetting function parameters or forgetting function form: Reply to Loftus and Bamber and to Wixted. *J. Exp. Psychol. Learn. Mem. Cogn.* 16: 936–945.

Bower GH (1981) Mood and memory. *Am. Psychol.* 36: 129–148.

Bower GH and Forgas JP (2000) Affect, memory, and social cognition. In: Forgas J, Eick K, and Bower GH (eds.) *Cognition and Emotion*, pp. 87–168. New York: Oxford University Press.

Brown GD, Preece T, and Hulme C (2000) Oscillator-based memory for serial order. *Psychol. Rev.* 107: 127–181.

Chechile RA (1987) Trace susceptibility theory. *J. Exp. Psychol. Gen.* 116: 203–222.

Chechile RA (2006) Memory hazard functions: A vehicle for theory development and test. *Psychol. Rev.* 113: 31–56.

Christensen H, Kopelman M, Stanhope N, Lorentz L, and Owen P (1998) Rates of forgetting in Alzheimer dementia. *Neuropsychologia* 36: 547–557.

Cowan N, Saults J, and Nugent LD (1997) The role of absolute and relative amounts of time in forgetting within immediate memory: The case of tone-pitch comparisons. *Psychon. Bull. Rev.* 4: 393–397.
Cowan N, Saults S, and Nugent L (2001) The ravages of absolute and relative amounts of time on memory. In: Roediger HL III, Nairne JS, Neath I, and Surprenant AM (eds.) *The Nature of Remembering: Essays in Honor of Robert G. Crowder*, pp. 315–330. Washington, DC: American Psychological Association.
Cowan N, Beschin N, and Della Sala S (2004) Verbal recall in amnesiacs under conditions of diminished retroactive interference. *Brain* 127: 825–834.
Dudai Y (2004) The neurobiology of consolidations, or, how stable is the engram. *Annu. Rev. Psychol.* 55: 51–86.
Ebbinghaus H (1885/1964) *Memory: A Contribution to Experimental Psychology*. Oxford: Dover.
Estes W (1955) Statistical theory of spontaneous recovery and regression. *Psychol. Rev.* 62: 145–154.
Freud S (1915) *Repression*, vol. XIV. London: Hogarth.
Galton F (1879) Psychometric experiments. *Brain* 2: 148–162.
Gleaves DH, Smith SM, Butler LD, and Spiegel D (2004) False and recovered memories in the laboratory and clinic: A review of experimental and clinical evidence. *Clin. Psychol. Sci. Pract.* 11: 3–28.
Greene RL (1996) The influence of experimental design: The example of the Brown-Peterson paradigm. *Can. J. Exp. Psychol.* 50: 240–242.
Jenkins J and Dallenbach K (1924) Obliviscence during sleep and waking. *Am. J. Psychol.* 35: 605–612.
Jost A (1897) Die Assoziationsfestigkeit in ihrer Abhängigkeit von der Verteilung der Wiederholungen [The strength of associations in their dependence on the distribution of repetitions]. *Z. Psychol. Physiol. Sinnesorgane* 14: 436–472.
Kareev Y (2000) Seven (indeed, plus or minus two) and the detection of correlations. *Psychol. Rev.* 107: 397–403.
Keppel G (1968) Retroactive and proactive inhibition. In: Dixon TR and Horton DL (eds.) *Verbal Behavior and General Behavior Theory*, pp. 172–213. Englewood Cliffs, NJ: Prentice-Hall.
Kihlstrom JF (2004) An unbalanced balancing act: Blocked, recovered, and false memories in the laboratory and clinic. *Clin. Psychol. Sci. Pract.* 11: 34–41.
Koppenaal R (1963) Time changes in strength of A-B, A-C lists: Spontaneous recovery? *J. Verb. Learn. Verb. Behav.* 2: 310–319.
Lee MD (2004) A Bayesian analysis of retention functions. *J. Math. Psychol.* 48: 310–321.
Levy BJ and Anderson MC (2002) Inhibitory processes and the control of memory retrieval. *Trends Cogn. Sci.* 6: 299–305.
Loftus GR and Bamber D (1990) Learning-forgetting independence, unidimensional memory models, and feature models: Comment on Bogartz (1990). *J. Exp. Psychol. Learn. Mem. Cogn.* 16: 916–926.
Luria AR (1968) *The Mind of a Mnemonist*. New York: Basic Books.
MacDonald SW, Stigsdotter-Neely A, Derwinger A, and Backman L (2006) Rate of acquisition, adult age, and basic cognitive abilities predict forgetting: New views on a classic problem. *J. Exp. Psychol. Gen.* 135: 368–390.
MacLeod MD and Macrae C (2001) Gone but not forgotten: The transient nature of retrieval-induced forgetting. *Psychol. Sci.* 12: 148–152.
McBride DM and Dosher BA (1999) Forgetting rates are comparable in conscious and automatic memory: A process-dissociation study. *J. Exp. Psychol. Learn. Mem. Cogn.* 25: 583–607.
McGeoch JA (1932) Forgetting and the law of disuse. *Psychol. Rev.* 39: 352–370.
McGeoch JA (1942) *The Psychology of Human Learning*. New York: Longmans and Green.
McGeoch JA and Irion AL (1952) *The Psychology of Human Learning*, 2nd edn. New York: Longmans and Green.
Meeter M and Murre JM (2004) Consolidation of long-term memory: Evidence and alternatives. *Psychol. Bull.* 130: 843–857.
Melton AW and Irwin JM (1940) The influence of degree of interpolated learning on retroactive inhibition and the overt transfer of specific responses. *Am. J. Psychol.* 53: 173–203.
Mensink G-J and Raaijmakers JG (1988) A model for interference and forgetting. *Psychol. Rev.* 95: 434–455.
Müller GE and Pilzecker A (1900) Experimentelle Beiträge zur Lehre vom Gedächtnis [Experimental contributions to the science of memory]. *Z. Psychol.* 1(supplement): 1–300.
Nairne JS (2001) A functional analysis of primary memory. In: Roediger HL III, Nairne JS, Neath I, and Surprenant AM (eds.) *The Nature of Remembering: Essays in Honor of Robert G. Crowder*, pp. 283–296. Washington, DC: American Psychological Association.
Nairne JS (2002a) The myth of the encoding-retrieval match. *Memory* 10: 389–395.
Nairne JS (2002b) Remembering over the short-term: The case against the standard model. *Annu. Rev. Psychol.* 53: 53–81.
Osgood CE (1953) *Method and Theory in Experimental Psychology*. New York: Oxford University Press.
Page MP and Norris D (1998) The primacy model: A new model of immediate serial recall. *Psychol. Rev.* 105: 761–781.
Pavlov IP (1927) *Conditioned Reflexes*. London: Oxford University Press.
Pitt MA, Myung IJ, and Zhang S (2002) Toward a method of selecting among computational models of cognition. *Psychol. Rev.* 109: 472–491.
Porter S, Yuille JC, and Lehman DR (1999) The nature of real, implanted, and fabricated memories for emotional childhood events: Implications for the recovered memory debate. *Law Hum. Behav.* 23: 517–537.
Postman L (1961) The present status of interference theory. In: Cofer CN and Musgrave BS (eds.) *Verbal Learning and Verbal Behavior*, pp. 152–196. New York: McGraw-Hill.
Postman L and Underwood BJ (1973) Critical issues in interference theory. *Mem. Cognit.* 1: 19–40.
Reitman JS (1974) Without surreptitious rehearsal, information in short-term memory decays. *J. Verb. Learn. Verb. Behav.* 13: 365–377.
Ribot T (1881) *Les maladies de la mémoire [Diseases of Memory]*. Paris: Germer Bailliere.
Roberts S and Pashler H (2000) How persuasive is a good fit? A comment on theory testing. *Psychol. Rev.* 107: 358–367.
Roediger HL III (1973) Inhibition in recall from cueing with recall targets. *J. Verb. Learn. Verb. Behav.* 12: 644–657.
Roediger HL III (1974) Inhibiting effects of recall. *Mem. Cognit.* 2: 261–269.
Roediger HL III (1978) Recall as a self-limiting process. *Mem. Cognit.* 6: 54–63.
Roediger HL III and Challis BH (1989) Hypermnesia: Increased recall with repeated tests In: Izawa C (ed.) *Current Issues in Cognitive Processes: The Tulane Flowerree Symposium on Cognition*, pp. 175–199. Hillsdale, NJ: Erlbaum.
Roediger HL III, Knight JL, and Kantowitz BH (1977) Inferring decay in short-term memory: The issue of capacity. *Mem. Cognit.* 5: 167–176.
Rubin DC, Hinton S, and Wenzel AE (1999) The precise time course of retention. *J. Exp. Psychol. Learn. Mem. Cogn.* 25: 1161–1176.
Rubin DC and Wenzel AE (1996) One hundred years of forgetting: A quantitative description of retention. *Psychol. Rev.* 103: 734–760.

Salthouse TA (1991) *Theoretical Perspectives on Cognitive Aging*. Hillsdale, NJ: Erlbaum.
Schacter DL and Addis DR (2007) The cognitive neuroscience of constructive memory: Remembering the past and imagining the future. *Philos. Trans. R. Soc. B* 362: 773–786.
Schooler LJ and Hertwig R (2005) How forgetting aids heuristic inference. *Psychol. Rev.* 112: 610–628.
Schooler JW, Ambadar Z, and Bendiksen M (1997) A cognitive corroborative case study approach for investigating discovered memories of sexual abuse. In: Read JD and Lindsay DS (eds.) *Recollections of Trauma: Scientific Evidence and Clinical Practice*, pp. 379–387. New York: Plenum Press.
Shobe KK and Kihlstrom JF (2002) Interrogative suggestibility and ''memory work.'' In: Eisen ML, Quas JA, and Goodman GS (eds.) *Memory and Suggestibility in the Forensic Interview*, pp. 309–327. Mahwah, NJ: Erlbaum.
Simon HA (1966) A note on Jost's law and exponential forgetting. *Psychometrika* 31: 505–506.
Slamecka NJ (1985) On comparing rates of forgetting: Comment on Loftus (1985) *J. Exp. Psychol. Learn. Mem. Cogn.* 11: 812–816.
Slamecka NJ and McElree B (1983) Normal forgetting of verbal lists as a function of their degree of learning. *J. Exp. Psychol. Learn. Mem. Cogn.* 9: 384–397.
Suddendorf T and Corballis MC (1997) Mental time travel and the evolution of the human mind. *Genet. Soc. Gen. Psychol. Monogr.* 123: 133–167.
Tulving E (1974) Cue-dependent forgetting. *Am. Sci.* 62: 74–82.
Tulving E (1983) *Elements of Episodic Memory*. New York: Oxford University Press.
Tulving E (2002) Chronesthesia: Conscious awareness of subjective time. In: Stuss DT and Knight RT (eds.) *Principles of Frontal Lobe Function*, pp. 311–325. New York: Oxford University Press.
Tulving E and Pearlstone Z (1966) Availability versus accessibility of information in memory for words. *J. Verb. Learn. Verb. Behav.* 5: 381–391.
Tulving E and Thomson DM (1973) Encoding specificity and retrieval processes in episodic memory. *Psychol. Rev.* 80: 352–373.
Turvey M, Brick P, and Osborn J (1970) Proactive interference in short-term memory as a function of prior-item retention interval. *Q. J. Exp. Psychol.* 22: 142–147.
Underwood BJ (1957) Interference and forgetting. *Psychol. Rev.* 64: 49–60.
Underwood BJ (1964) Degree of learning and the measurement of forgetting. *J. Verb. Learn. Verb. Behav.* 3: 112–129.
Watkins OC and Watkins MJ (1975) Buildup of proactive inhibition as a cue-overload effect. *J. Exp. Psychol. Hum. Learn. Mem.* 1: 442–452.
Waugh NC and Norman DA (1965) Primary memory. *Psychol. Rev.* 72: 89–104.
Wheeler MA (1995) Improvement in recall over time without repeated testing: Spontaneous recovery revisited. *J. Exp. Psychol. Learn. Mem. Cogn.* 21: 173–184.
Wickelgren WA (1977) *Learning and Memory*. Englewood Cliffs, NJ: Prentice-Hall.
Williams LM (1994) Recall of childhood trauma: A prospective study of women's memories of child sexual abuse. *J. Consult. Clin. Psychol.* 62: 1167–1176.
Wixted JT (2004a) On common ground: Jost's (1897) law of forgetting and Ribot's (1881) law of retrograde amnesia. *Psychol. Rev.* 111: 864–879.
Wixted JT (2004b) The psychology and neuroscience of forgetting. *Annu. Rev. Psychol.* 55: 235–269.
Wixted JT and Carpenter SK (2007) The Wickelgren power law and the Ebbinghaus savings function. *Psychol. Sci.* 18: 133–134.
Wixted JT and Ebbesen EB (1991) On the form of forgetting. *Psychol. Sci.* 2: 409–415.