7 Types and consequences of forgetting: Intended and unintended

Elizabeth L. Bjork, Robert A. Bjork, and Malcolm D. MacLeod

For most of us, remembering is good; forgetting is bad – something we wish to avoid. For most of us, in fact, "forgetting things" is the biggest complaint we have about our memories. Contrary to such intuitions, however, forgetting is a necessary and critical aspect of an efficient and adaptive memory system and, often, exactly what we must do to keep our memories functioning optimally. When previously stored information becomes out of date or no longer functional, we need some way to set aside, suppress, or forget that old information. Additionally, when we search our memories for some desired piece of information, such as the name of a former colleague or the source of an article that we have recently read, we need to inhibit or "forget" closely related, but incorrect, pieces of information that compete for retrieval with the target of our search. As we and others have previously argued (e.g., R. A. Bjork, 1989; E. L. Bjork & Bjork, 1996; E. L. Bjork, Bjork, & Anderson, 1998; Macrae & MacLeod, 1999), without such active forgetting mechanisms, we would soon become incapable of retrieving the information we need now, owing to interference from all the related information we have learned in the past.

In some cases, the types of forgetting that serve our goals and needs seem intentional, or at least not unintentional, as when we deliberately try to update our memories by forgetting or replacing an old phone number or an old password with a new one, or when we try to avoid retrieving uncomfortable or painful past events in our lives as a way of reducing their potentially aversive influences on our present lives. Such goal-directed forgetting may occur in response to explicit or implicit cues to forget, initiated either by ourselves or others, in both real-world and laboratory situations. And, contrary to what our intuitions about our memories might suggest – that even should we wish to forget certain previously stored information in response to an explicit cue, we would not be able to do so – extensive research using a procedure known as the intentional- or directed-forgetting paradigm (e.g., R. A. Bjork, 1972, 1989) shows that individuals can, in fact, forget previously stored information when instructed to do so.

In other types of forgetting important for the efficient functioning of our memories, however, the forgetting is unintended, may not be consistent with

our long-term goals, and is instigated by cues of a more subtle and implicit nature. In the process of updating our memories so as to remember, for example, where we parked our car this morning, rather than yesterday morning, we typically do not give ourselves an explicit cue to forget the previous parking spot; rather, the cue to forget is implicit, and the forgetting or suppression of the old parking spot most likely occurs as a consequence of our noting and retrieving to ourselves the current parking spot. Analogously, in laboratory research on this type of forgetting using the retrieval-practice procedure of Anderson, Bjork, and Bjork (1994), the instruction to forget is not explicit, as it is in the intentional- or directed-forgetting paradigm; rather, it arises as an implicit and subtle component of the procedure itself. As participants practice retrieving some of the items associated with a particular cue during a retrieval-practice phase, there is an implicit instruction to forget, inhibit, or suppress the other items that were associated with that same cue during a preceding study phase and are now competing with the desired target of retrieval. In selecting for the desired target and against these competitors, the latter become suppressed or forgotten. Although such retrieval-induced forgetting serves the immediate goal of the participant – successful retrieval of the desired target during the retrieval-practice phase – it is difficult to conceive of such forgetting as being consciously intentional, nor does it serve the participant's long-term goal, which is to remember as many of the cue–target pairs presented during the original study episode as possible.

Whether occurring intentionally or unintentionally, however, or in response to an explicit or implicit cue, we believe that these types of forgetting are critical to the efficient functioning of human memory and are driven by a common underlying inhibitory process or mechanism, which we and others have called retrieval inhibition (e.g., R. A. Bjork, 1989; E. L. Bjork, Bjork, & Anderson, 1998; Geiselman, Bjork, & Fishman, 1983; Macrae & MacLeod, 1999; Roediger, 1974), and by which we mean the loss of retrieval access to information that nonetheless remains available in memory. That is, although our retrieval access to such inhibited information has become impaired or diminished, the inhibited or suppressed information still remains stored and available in memory, as can be demonstrated by other measures, such as recognition tests.

In the remainder of this chapter, we first briefly describe the principal paradigms used to study intentional forgetting, the basic characteristics of such forgetting, and the role of retrieval inhibition in its production. We then discuss more fully the type of forgetting that is our primary focus – retrieval-induced forgetting – including a description of the retrieval-practice paradigm recently developed to study it, the basic characteristics of such forgetting that implicate the role of inhibition or suppression in its production, and certain recently discovered extensions and boundary conditions of retrieval-induced forgetting. Finally, we explore several areas in which retrieval-induced forgetting may have unintended, but nonetheless profound,

consequences in our lives: reliability of witness memory, stereotyping and other forms of impression formation, and – possibly – repression.

Intentional forgetting: Basic paradigms, results, and the role of retrieval inhibition

In perhaps the most typical instantiation of the paradigm used to study intentional forgetting, following the presentation of a list of words (List 1) that is to be learned for a later memory test, some participants are then directed "to forget" that list, which they are told was presented just for practice or by mistake, and are then presented with the "real" or the "correct" list to learn (List 2). In contrast, a different group of participants are presented with List 1, told to keep remembering it, and are then presented with List 2 to learn as well. In another variation of this paradigm, the instruction "to forget" does not come as a surprise; rather, all participants are informed at the beginning of the study that half way through each list of words presented, a cue either to forget or to remember the preceding words (i.e., List 1) will occur. Then, immediately following this mid-list cue, the remaining words, which are always to be remembered, are presented (i.e., List 2). Occasionally, a control condition is also employed in which the pre-cue or List-1 words are replaced by non-verbal materials to be processed in some way, and then the List-2 words are presented for study.

Across many studies employing such directed-forgetting procedures, a basic pattern of results has emerged that can be summarized in terms of three robust findings. First, List-2 words are significantly better recalled when participants are instructed to forget the preceding List-1 words than when they are told to keep remembering them; and, second, if a control condition was employed, List-2 words are frequently remembered just as well when participants were instructed to forget the List-1 words as when no words were presented at all in List 1. Third, if participants are unexpectedly asked to recall words that they were instructed to forget, their ability to do so is significantly impaired compared to their ability to recall words that they were instructed to remember. On the basis of the first two findings, it appears that directing people to forget previously studied material significantly reduces, and can even eliminate, the proactive interference of the to-be-forgotten material on subsequent learning. Based on the third finding, directing people to forget previously studied material appears to impede their later conscious access to that material or to the original learning episode that is the target of the forget instruction. (For a review of the studies from which this basic pattern of intentional-forgetting results has been extrapolated, see R. A. Bjork, 1972, 1989; Johnson, 1994; and C. M, MacLeod, 1998.)

As previously indicated, we believe that the most compelling explanation for this pattern of directed-forgetting results is in terms of retrieval inhibition. Namely, when individuals are instructed to forget previously studied information and then given new information to learn in its place, a process is

initiated that inhibits subsequent retrieval access to the to-be-forgotten information or the learning episode that was the object of the instruction to forget. Moreover, because the to-be-forgotten information is not retrievable, it does not interfere with the recall of the current to-be-remembered information; thus, the decrease or elimination of proactive interference owing to the to-be-forgotten information as well as its significantly diminished recall relative to that for comparable to-be-remembered information.

Furthermore, although one's retrieval access to previously studied information becomes inhibited as a consequence of being instructed to forget it, the strength in memory of such information appears to be left unaffected, as indicated by the following three findings. First, recognition memory for the to-be-forgotten information is unimpaired as compared to that for comparable to-be-remembered information (e.g., Block, 1971; Elmes, Adams, & Roediger, 1970; Geiselman et al., 1983). Second, in a relearning paradigm, to-be-forgotten items can be learned as readily as to-be-remembered items (e.g., Reed, 1970; Geiselman & Bagheri, 1985). Third, under certain conditions, the proactive interference owing to the forgotten list can be reinstated at full strength, such as when some to-be-forgotten items are encountered as foils on a recognition test for to-be-remembered items (E. L. Bjork, Bjork, & Glenberg, 1973; E. L. Bjork & Bjork, 1996). Finally, observations based on indirect measures of retention provide evidence for continuing indirect or implicit influences of information that has been intentionally forgotten, even when retrieval access to such information remains inhibited (e.g., Basden, Basden, & Gargano, 1993; E. L. Bjork, Bjork, & Kilpatrick, 1990; E. L. Bjork & Bjork, 1996). Indeed, under certain circumstances, the indirect influences of previously studied information appear to be larger when that information has been intentionally forgotten versus intentionally remembered (E. L. Bjork & Bjork, 2003; E. L. Bjork, R. A. Bjork, Stallings, & Kimball, 1996).

Thus, in terms of the present analysis, it seems reasonable to think of the type of forgetting observed in response to an instruction to forget – whether implicit or explicit and whether administered by ourselves in everyday life or by an experimenter in the laboratory – to have two intended consequences of a positive nature, and one unintended consequence that, under certain circumstances at least, could be negative. First, because the process so initiated inhibits the subsequent retrieval of the information that is the target of the forget instruction, such information does not interfere with our learning and recall of the new information that is to replace it, thereby allowing for efficient memory updating and ease of access to the new information. Second, because the representation of the intentionally forgotten information in our memories appears to remain intact and at essentially full strength, as indicated by other measures such as recognition and priming, should our circumstances change and we need to regain access to such information, we can become fluent in its use again more quickly than were we to have to learn it anew from scratch. On perhaps the negative side, however, because such intentionally forgotten information, although still inhibited, can continue to

influence our behaviour in indirect ways, there can be unintended consequences or influences of such forgetting about which we are unaware and thus cannot consciously correct for or mitigate.

Retrieval-induced forgetting: Basic paradigms, results, and the role of retrieval inhibition

Although a considerable body of results – obtained over many years of research employing a variety of paradigms – supports the notion that the act of retrieving some information impairs one's ability to retrieve other related information, we focus in the present chapter on the type of forgetting that has come to be called retrieval-induced forgetting as demonstrated in the retrieval-practice procedure recently developed by Anderson et al. (1994).

The standard version of this procedure employs four distinct phases: a study phase, a directed retrieval-practice phase, a distractor phase, and a final test phase. In the study phase, participants are presented with a list of category-exemplar pairs to learn (e.g., *Fruit Orange; Weapon Rifle*) – typically, six such pairs from each of eight categories, presented individually and in a mixed order. Next, participants practice retrieving half of the exemplars from half of the studied categories (e.g., *Fruit*), and none of the exemplars from the remaining categories (e.g., *Weapon*). This retrieval practice is directed in that a category name together with a letter stem (e.g., *Fruit Or* ____) is presented with participants instructed to retrieve the previously studied exemplar that fits the combined category-stem cue. Also, to increase the effectiveness of this retrieval practice, each practiced pair is typically given three such tests separated by expanding intervals filled with the retrieval practice of other pairs (Landaur & Bjork, 1978). Then, following the distractor phase, originally lasting 20 minutes, the final and unexpected recall test is administered by presenting each category cue (e.g., *Fruit, Weapon*) and asking participants to recall all the members of that category that they can remember from any part of the experiment.

On this final test, the recall of three types of items – defined by the preceding study and retrieval-practice phases – is of interest: (1) practiced exemplars from practiced categories (e.g., *Orange*; called Rp+ items); (2) unpracticed items from practiced categories (e.g., *Banana*, called Rp– items); and (3) unpracticed items from unpracticed categories, (e.g., *Rifle*; called Nrp or baseline items). In studies using this procedure, performance on the final test typically reveals the following pattern: whereas, the final recall of items that received retrieval practice, the Rp+ items, is significantly facilitated relative to recall of the baseline Nrp items, the final recall of the unpracticed items from the practiced categories, the Rp– items, is significantly impaired relative to this same baseline. It is also important to note that when final-test performance is assessed using a cued-recall procedure in which recall of non-practiced Rp– items is required before recall of Rp+ items within a category – thus controlling for output interference or the possibility that the observed impaired

recall of Rp– items is entirely due to the earlier recall of the stronger Rp+ items during the final test – this pattern of results is still observed (e.g., Anderson et al., 1994, Experiment 2). Additionally, post-hoc analyses of free-recall performance (e.g., Macrae & MacLeod, 1999) have also precluded output interference as a significant factor in the production of the impaired recall of Rp– items.

This pattern of results, which has been replicated across a number of studies using both verbal materials such as those described above (Anderson et al., 1994; Anderson & Spellman, 1995) and visuospatial materials (Ciranni & Shimamura, 1999) is consistent with the notion that the act of retrieving some items impairs the retrieval of related items. Because the resultant impairment appears to be initiated during the retrieval-practice phase and can last at least 20 minutes, this pattern has come to be called retrieval-induced forgetting. Furthermore, as previously argued (Anderson et al., 1994; E. L. Bjork et al., 1998; Macae & MacLeod, 1999), we believe that inhibitory processes are responsible for the impaired recall of the unpracticed, related items. More specifically, our conception of the mechanism producing such inhibition is as follows: During the retrieval practice of *Orange* to the cue Fruit, other exemplars, such as *Banana*, are also activated, creating competition and interference. To retrieve the desired target, *Orange*, in the presence of this competition, *Banana* must be selected against or inhibited, with the resultant suppression of *Banana* leading to its impaired recall on the final test.

It needs to be pointed out, however, that – although our interpretation of these basic retrieval-induced forgetting results is in terms of inhibition or suppression – other explanations not involving the assumption of inhibition, such as blocking (e.g., Blaxton & Neely, 1983; Roediger, 1974; Roediger & Neely, 1982; Rundus, 1973) or response competition owing to the strengthening of the practiced exemplars (e.g., McGeoch, 1942; Mensink & Raajimakers, 1988; Raajimakers & Shiffrin, 1981) are also consistent with this basic pattern. We believe, however, that additional studies designed to test among these proposed explanations favour an account in terms of inhibition. In these studies, the research strategy has been to construct a situation in which unpracticed exemplars would not be expected to compete – and thus would not need to be selected against or inhibited – during the retrieval practice of their related category-exemplar pairs. Then, having created such a situation, to see whether the later recall of such unpracticed exemplars is nonetheless impaired, which would be consistent with blocking, response-competition, or strength-dependent explanations, or whether their later recall is not impaired, which would be consistent with the suppression explanation.

To illustrate, in one such study employing this strategy, Anderson et al. (1994) manipulated the taxonomic strength of practiced and unpracticed exemplars, with the assumption that taxonomically strong, unpracticed exemplars would compete during retrieval practice of other category-exemplar pairs and, thus, have to be suppressed; whereas, taxonomically weak,

unpracticed exemplars would be less likely to compete during retrieval practice of other category-exemplar pairs and, thus, escape being suppressed. In contrast, most non-inhibitory based explanations would have to predict that the later recall of either type of unpracticed exemplar would be impaired owing to the strengthening of the related practiced pairs. The pattern of results obtained supported the inhibitory-based explanation. Recall of unpracticed taxonomically strong exemplars (i.e., the ones that should compete and thus need to be suppressed during retrieval practice for related exemplars) was impaired on the final test, whether their related practiced exemplars were taxonomically strong (e.g., *Orange*) or taxonomically weak (e.g., *Papaya*); whereas, recall of unpracticed taxonomically weak exemplars (i.e., the ones that should not compete and thus escape being suppressed during retrieval practice for related exemplars) was not impaired, whether their related practiced exemplars were taxonomically strong or weak. Furthermore, this pattern was obtained in the presence of large facilitative (or strengthening) effects of retrieval practice for both strong and weak exemplars.

Additionally, in a study manipulating the type of retrieval practice (Anderson, Bjork, & Bjork, 2000), results consistent with the inhibition explanation of retrieval-induced forgetting were likewise obtained. In this study, participants engaged in one of two types of retrieval-practice conditions: (1) competitive, which followed the same procedures of Anderson et al. (1994); and (2) non-competitive, in which, rather than recalling an exemplar in response to a category label and a two-letter stem (e.g., *Fruit Or* ____), participants were given the exemplar along with a two-letter stem for the category (e.g., *Fr* ____ *Orange*) and asked to recall the category label. Thus, in both conditions, participants practiced retrieving the same category-exemplar associations for the same number of times and, in both conditions, it was expected that those associations would be strengthened. Similar to the preceding study, however, it was assumed that only in the competitive condition would attempts to retrieve the target exemplar, *Orange*, invoke competition from other strong exemplars, such as *Banana*, which would then need to be suppressed. In contrast, in the non-competitive condition, no such competition among exemplars should arise and, thus, there should be no consequent suppression of the unpracticed exemplars, such as *Banana*. As expected, both types of practice resulted in significant strengthening of the practiced category-exemplar pairs; that is, recall of *Orange*, for example, was facilitated on the later recall test by both the *Fruit Or* ____ or *Fr* ____ *Orange* types of practice. Critically, however, recall of the unpracticed exemplars, such as *Banana*, was only impaired in the competitive condition, lending strong support to the assumption that inhibitory processes – invoked to resolve interference from competitors during attempts to recall a target item – underlie retrieval-induced forgetting.

Finally, perhaps the strongest support for an explanation of retrieval-induced forgetting in terms of inhibitory mechanisms comes from the

research of Anderson and Spellman (1995) who reasoned that if competing responses were suppressed in the attempt to retrieve the designated target during retrieval practice, then those inhibited items should be more difficult to retrieve, not only from the studied retrieval cue, but from any cue used to test them. Consistent with this assumption, Anderson and Spellman found that when participants practiced retrieving some exemplars of a category (e.g., *Red Blood*), their delayed recall of other unpracticed exemplars of that category was impaired, whether those exemplars were tested with the same cue used for practice (e.g., *Red*) or a different cue not used during practice (e.g., *Food*). The critical aspect of this finding in support of an inhibitory explanation is that the ability of participants to recall such items in response to the unpracticed category cue was impaired even though that cue was unrelated to the item strengthened during retrieval practice. In this situation, then, the unpracticed cue provides a measure of the accessibility of these related, unpracticed items that is independent of associative interference from the practiced targets. As argued by Anderson and Spellman, such cue-independent impairment provides strong support for the assumption of inhibitory processes acting during retrieval to suppress competing responses in order to focus retrieval search on the desired target.

Extensions of retrieval-induced forgetting

As previously indicated, we believe retrieval-induced forgetting to be prevalent in our everyday lives, playing a critical role in our frequent need to update our memories and to resolve interference during retrieval. Our making such a claim, however, immediately raises the issue of the generality of the phenomenon of retrieval-induced forgetting and whether this type of forgetting does, in fact, extend beyond the rather constrained learning and retrieval-practice procedures represented in the standard Anderson et al. (1994) paradigm. Several relevant questions thus come to mind. For example, would such forgetting occur when we are encoding information of a more personal or social relevance to us than the learning of category-exemplar pairs, and would it occur when we are not specifically encoding or processing such information with the conscious intention of learning it for a later memory test? Conversely, would it occur if, during retrieval practice for *some* of the items, it were fully understood that there would be a later recall test for *all* of the items. In the typical instantiation of the Anderson et al. retrieval-practice paradigm, the final recall test comes as a surprise, leaving open the question of whether inhibition of unpracticed competitors would occur were participants to know in advance that a later recall test for all of the originally studied category-exemplar pairs would be given. And, lastly, would retrieval-induced forgetting occur with less retrieval practice than typically provided? If the inhibitory processes assumed to underlie retrieval-induced forgetting play the important role in memory updating and the resolution of competition in everyday memory retrieval that we are suggesting, then it seems

unreasonable that such extensive retrieval practice would be a necessary condition for their elicitation.

These questions concerning the generality of retrieval-induced forgetting and the variety of conditions under which it occurs were recently raised and addressed by Macrae and MacLeod (1999) in a series of three studies. To address the first question, they presented participants with the social-cognition task of impression formation. Specifically, participants were asked to form impressions about two individuals, named *John* and *Bill*, on the basis of traits (e.g., *trustful, cheerful*), which were presented in a similar manner to that used to present the category-exemplar pairs in a typical Anderson et al. (1994) paradigm. Participants first looked through a stack of randomly ordered cards on each of which one name (i.e., *John* or *Bill*) was printed together with one of 10 traits that had been randomly paired with that name. Next, participants were presented with cue cards that probed their memory for half of the traits paired with one of the individuals, say *John*, making those traits the Rp+ items, while the remaining traits for *John* thus become Rp– items, and the traits for the other individual, *Bill*, become Nrp items. To illustrate, in this case, such a cue card would contain the name *John* along with a two-letter stem for a trait (e.g., *John tr*____ as a probe for retrieving *trustful*), with each such card presented on three occasions. Then, after a distractor phase, a surprise test was given in which participants were asked to recall all the traits they could remember that had described *John* and *Bill*, respectively. As in the intentional-learning situation of the Anderson et al. paradigm, the typical pattern of retrieval-induced forgetting was observed. Specifically, recall of the Rp+ traits was significantly facilitated compared to the Nrp or baseline traits, whereas recall of the Rp– traits was significantly impaired compared to that for the Nrp traits, demonstrating the occurrence of retrieval-induced forgetting in the context of the social-information processing task of impression formation and, furthermore, its generalization to a situation in which the original encoding of information did not take place in the context of an intentional leaning set.

To address the second question – whether retrieval-induced forgetting would emerge when participants knew in advance that a final recall test would be given and are thus motivated to retain all of the originally studied information – Macrae and MacLeod (1999) attempted to create a type of situation with which students are often faced: namely, needing to prepare for an examination that will cover all of the assigned material, but with only enough time to practice retrieving some of that material before the exam. To set up such a situation, participants were instructed that they would be presented with information to learn about two tropical islands (*Tok* and *Bilu*) in preparation for a later geography exam, and they were then given 10 facts to learn about each island (e.g., *The official language in Tok is French*). Next, participants engaged in a retrieval-practice phase in which they were presented with cue cards probing their memory for half the facts

associated with one of the islands, say *Bilu*, with each card containing the name of the island and a hint about the fact to be recalled (e.g., for the fact, "*Bilu's only major export is copper*", participants were cued with the hint, "*Bilu's only major export is c* ____") and with each such card presented on three occasions. Then, after a distractor task, participants were given a simulated final exam in which they were asked to recall as much as they could about both of the tropical islands. As in their previous study using names and traits, the pattern of retrieval-induced forgetting was observed: participants recalled significantly more practiced facts (Rp+ items) about an island than unpracticed facts about the different island (Nrp baseline items); whereas, compared to the same baseline, they recalled significantly fewer of the unpracticed facts about the practiced island (the Rp– items). These results thus extended the occurrence of retrieval-induced forgetting to situations in which participants, before the study phase, have foreknowledge of a final test and, thus, should be highly motivated not to forget any of the originally studied information.

Lastly, to address the question of the amount of retrieval practice required to produce retrieval-induced forgetting, Macrae and MacLeod (1999) varied the number of retrieval-practice trials given to different groups of participants in the context of the previously described impression-formation task. Specifically, the Rp+ items were presented for practice on one, three, or six occasions, with the duration of each practice trial timed in a manner to equate the duration of the total retrieval-practice phase for each group. The question of primary interest, of course, was how many practices would be necessary to elicit retrieval-induced forgetting. Consistent with the notion that such forgetting plays a critical role in the regulation of everyday memory, impaired retrieval of Rp– items compared to Nrp items was obtained in all three retrieval-practice conditions, establishing that a single retrieval practice of the Rp+ items can be sufficient to produce retrieval-induced forgetting for the related, non-practiced items.

Taken together, then, the results from these three studies lend compelling support to the contention that retrieval-induced forgetting is an adaptive and pervasive mechanism of memory, playing a critical role in the type of memory updating and the resolution of interference in retrieval that occurs in everyday life. On the other hand, retrieval-induced forgetting is not all pervasive. For example, as we have already seen in the studies manipulating the taxonomic strength of category-exemplar pairs (Anderson et al., 1994) and the nature of retrieval practice (Anderson et al., 2000), there are circumstances under which successful retrieval of desired targets occurs without producing impaired retrieval of associated material in memory. Moreover, as discussed in the next section, recent research from a number of investigators has revealed additional moderating factors or boundary conditions on the occurrence of retrieval-induced forgetting.

Boundary conditions on retrieval-induced forgetting

In research by Anderson and McCulloch (1999), retrieval-induced forgetting has been shown to be significantly reduced, and even eliminated, when participants are instructed to interrelate or integrate the different exemplars within a given category during the initial study phase. Although how integration acts to limit the occurrence of retrieval-induced forgetting is not yet fully understood, the results obtained by these investigators across a series of experiments point to diminished suppression and mediated retrieval as the two most likely explanations. The notion of diminished suppression as an explanatory factor builds upon the finding from numerous studies on fact retrieval that integration reduces retrieval competition (e.g., Moeser, 1979a, 1979b; E. E. Smith, Adams, & Schorr, 1978). If, similarly, linking items associated to the same cue via integration reduces retrieval competition among them, then there should be little need to suppress non-target items during retrieval practice of target items. The notion of mediated retrieval as an explanatory factor stems from studies conducted during the classical interference era of verbal learning demonstrating that participants could use connections established between competing associates as indirect retrieval routes, thereby evading both proactive and retroactive interference (for reviews, see Anderson & Neely, 1996; Postman, 1971). As applied to the phenomenon of retrieval-induced forgetting, the notion would be that links established between exemplars of the same category during study could then, during the final recall test, be used as indirect retrieval routes to recover suppressed non-target items. Suppose, for example, that an *Orange–Banana* link had been established via integration during study. Now, even if *Banana* is suppressed during the retrieval practice of *Orange* to the cue *Fruit*, such practice should also have enhanced the accessibility of the *Orange–Banana* link, allowing *Banana* to be retrieved via this link even if still suppressed.

In the first case, integration could be thought of as insulating non-target items from the suppression that produces retrieval-induced forgetting; in the latter case, integration could be thought of as offsetting or masking the suppression that normally produces retrieval-induced forgetting. Indeed, as Anderson and McCulloch (1999) conclude, the benefits of integration in reducing retrieval-induced forgetting most likely arise from a combination of reduced suppression and mediated retrieval. Importantly, as also argued by Anderson and McCulloch, these limiting factors on the occurrence of retrieval-induced forgetting may demonstrate the means by which the complex knowledge structures of experts are spared such forgetting, even though experts may repeatedly engage in the access of some, but not all, of the knowledge accumulated within their domains of expertise.

In addition to integration apparently protecting material from retrieval-induced forgetting, research by R. E. Smith and Hunt (2000) has demonstrated that distinctive encoding of information also seems to increase its resistance to such inhibition. To illustrate, when Smith and Hunt instructed

participants, during the initial study, to think about how each member of a category differed from all the other members of the category, subsequent retrieval practice given to a subset of the category members did not result in retrieval-induced forgetting of the remaining category members. (In contrast, it should be pointed out that when participants are instructed to encode, during initial study, similarities among items that are to become the competitors of the practiced targets, the retrieval inhibition suffered by such items can be significantly increased. For an in-depth discussion and theoretical explanation of how similarity relationships between practiced targets and competitors and among competitors themselves can lead to both increases and decreases in retrieval inhibition or suppression, see Anderson, Green, & McCulloch, 2000.)

The potential power of both distinctive encoding and integration to render information resistant to retrieval inhibition is illustrated in an interesting study by Macrae and Roseveare (2002). Using a variation of the standard retrieval-induced forgetting procedure, different participants were asked to learn a list of gifts by imagining them to have been purchased by (a) themselves, (b) their best friend, or (c) an unspecified other, with the expectation that participants receiving the first instruction would be led to engage in more self-relevant processing of the to-be-learned information than would the participants receiving the other two types of instruction. The intent of this manipulation was thus to test whether information encoded in terms of the self – presumed to constitute a highly complex knowledge structure (e.g., Kihlstrom & Cantor, 1984; Klein & Kihlstrom, 1986) – would spontaneously result in its being both more highly integrated and more distinctively processed, thus potentially protecting it from retrieval inhibition. As expected, whereas participants in the "best friend" and "other" instruction groups showed significant retrieval-induced forgetting on the final recall test, participants in the "self" instruction group did not.

As a check on this interpretation, Macrae and Roseveare (2002) also had participants, following the retrieval practice and final recall phases, rate the extent to which they had engaged in thoughts about potential recipients for the gifts, with the notion that such ratings would serve as a measure of the degree to which spontaneous distinctive processing in terms of the self had taken place during original study. These ratings confirmed that participants in the "self" condition had indeed engaged in more self-relevant and distinctive encoding of the gifts than had the other two groups of participants. Thus, it seems possible that the same distinctive (Smith & Hunt, 2000) and integrative (Anderson & McCulloch, 1999) encoding processes that have been shown to eliminate retrieval-induced forgetting in the laboratory may be spontaneously evoked whenever we engage in self-relevant processing of information, thereby setting another boundary condition on the occurrence of retrieval inhibition as produced via retrieval practice (see, M. D. MacLeod, Bjork, & Bjork, 2003, for a discussion of this possibility with respect to emotive material).

Finally and critically, it needs to be noted that there are temporal boundaries on the duration of retrieval inhibition as well. That is, the retrieval-induced forgetting of competing material in memory is temporary, resulting from decreases in its retrieval access, rather than its erasure or elimination from memory. (Or, as expressed in R. A. Bjork & Bjork's, 1992, new theory of disuse, resulting from decreases to its retrieval strength, not to its storage strength.) As we have also argued elsewhere (e.g., R. A. Bjork, 1989, E. L. Bjork et al., 1998; M. D. MacLeod et al., 2003; M. D. MacLeod & Macrae, 2001), this temporary nature of retrieval inhibition is adaptive. While in force, the inhibition serves to aid recall of desired information by suppressing interference from competing information. Should such inhibited information be needed in the future, however, it can become recallable again once its inhibition has been released or dissipated. Furthermore, as reasoned by M. D. MacLeod and Macrae (2001), if retrieval inhibition is to serve an adaptive memorial function, it must be transient in nature, owing to the constantly changing nature of our goal states, and, indeed, they have demonstrated that retrieval inhibition is no longer present following a 24-hour delay. If, however, a 24-hour delay is inserted between initial study and retrieval practice (rather than between retrieval practice and final recall), retrieval inhibition still occurs. Taken together, these two results demonstrate an interesting feature of retrieval inhibition – although transient in nature, it can be evoked for older memories, as well as newly established ones. It thus seems possible that selective retrieval practice over longer periods of time could well result in more permanent inhibition of competitive memories. (For further discussions of the potential relevance of repeated retrieval practice for such issues as the maintenance of one's self-image, the construction of autobiographical memories, and the suppression of unpleasant and/or traumatic memories, see E. L. Bjork et al., 1998, and M. D. MacLeod et al., 2003.)

Unintended consequences of retrieval-induced forgetting

As we have noted, an important aspect of the inhibition that we believe to underlie the phenomenon of retrieval-induced forgetting is that it occurs automatically or without our conscious attention. As we retrieve or remember some information, other related information is inhibited, with the inhibition occurring rather like a by-product. Hence, as we have previously suggested, there can be unintended consequences of such forgetting. For example, while the retrieval inhibition of competitors produced in the retrieval-practice phase of the Anderson et al. (1994) paradigm serves the immediate goal of successfully retrieving the desired target of the retrieval practice cue, it does not serve the longer-term goal of performing well on the final recall test for all the studied pairs. Moreover, as so aptly demonstrated in the study by Macrae and MacLeod (1999) in which participants learned geography facts for a later examination, this negative consequence of retrieval-induced forgetting occurs even when participants know in advance

that a final test for all the information studied will be given. Although motivated to remember all the originally studied facts in order to perform well on the expected final test, the retrieval practice given to a subset of the facts about a particular island nonetheless diminished the participants' ability to recall the remaining, unpracticed facts about that island. Worse yet, their recall for such unpracticed items was impaired relative to that of control participants who did not engage in any practice for any of the facts about either island. In the remainder of this section, several such unintended consequences of the phenomenon of retrieval-induced forgetting are discussed, some of them clearly negative, but others of perhaps a more positive nature.

Witness memory One of the more interesting areas in which the potential negative consequences of intentional-forgetting have been investigated is that of witness memory. In the first study to look at the possible relevance of retrieval-induced forgetting effects on eyewitness testimony, Shaw, Bjork, and Handal (1995) explored potential unintended consequences of the repeated questioning of witnesses on their long-term ability to recall details about the observed event. Although it had been realized for some time that repeated post-event interrogations of witnesses could alter their memories for the information that was the target of such questioning (e.g., Wells, 1993), Shaw et al. wondered – given the retrieval-induced forgetting observed by Anderson et al. (1994) – if such post-event interrogations might also have an impact upon the ability of witnesses to recall information that was *not* the object of such interrogations.

To explore this possibility, these investigators adapted the typical retrieval-practice procedure to make it suitable for an investigation of eyewitness memory. Participants were told to imagine that they had attended a party in a fellow student's apartment and, when leaving, noticed that their wallet was missing from their jacket pocket. They were then shown a series of 21 slides depicting objects in the bedroom where their jacket had been during the party, being told to pay close attention to the contents of the slides so that they would later be able to assist the police in an investigation of the theft. Within the slides presented were an assortment of objects that could typically be found in a student's bedroom, such as a desk, a computer, a telephone, and various personal items. Additionally, two categories of target items – college sweatshirts and college schoolbooks, with eight exemplars of each – were also contained within the slide sequence. After viewing the slides, participants were exposed to an "interrogation" phase in which they were given retrieval practice on half of the items in one of the two categories (e.g., four sweatshirts). To optimize the effects of this practice as is presumably achieved in the expanding retrieval procedure of Landauer and Bjork (1978), their interrogation phase included three retrieval-practice sessions that were (a) separated by increasingly longer intervals filled with distractor tasks and (b) designed to demand progressively more difficult retrievals. Thus, in the

first interrogation phase, participants were presented with fairly simple true–false questions about the four target items, asked from the standpoint of the owner of the apartment, *Janet*, and suggesting the correct response (e.g., "I think my friend Julia wore her Harvard sweatshirt. Was there a Harvard sweatshirt on my desk?"). In the second interrogation phase, participants were presented with more difficult cued-recall questions about the same target items asked from the perspective of a police officer who had arrived on the scene (e.g., "Was there a grey sweatshirt on Janet's desk? If so, what was the name of the university on that sweatshirt?"). In the third and final interrogation phase, the participants were presented with the most difficult cued-recall questions about the same items asked from the perspective of a detective investigating the case (e.g., "Were there any sweatshirts on the desk? What was written on those sweatshirts?"). Following this last interrogation phase and a distractor interval, participants engaged in a final "testimony phase", consisting of a category-cued-recall test in which they were asked to recall as many items in each of the two categories, sweatshirts and schoolbooks, as they could remember. Additionally, Shaw et al. included a group of participants who received no interrogation or retrieval practice. Instead, after viewing the slides, these no-interrogation control participants only engaged in distractor tasks until being given the same final recall test as the interrogated participants.

On the final recall test, the practiced or Rp+ items (e.g., the four sweatshirts that were the targets of the interrogation questions in the present example) were recalled at a significantly higher rate than were the Nrp items (e.g., the schoolbooks in the present example), whereas the Rp– items (e.g., the other four sweatshirts in the present example) were recalled at a significantly lower rate than were the Nrp items. Furthermore, because of the inclusion of a no-interrogation control condition, recall of the Rp+, Rp– and Nrp items by the interrogated participants could also be compared to the mean recall of participants who had not engaged in any retrieval practice for any of the items. Compared to the mean recall of this control group, the interrogated participants' recall of Rp+ items was significantly better, their recall of Rp– items was significantly worse, whereas their recall of Nrp items did not differ, indicating that Nrp items do not suffer from the retrieval practice given to a subset of items in another category. Prior to the inclusion of this no-interrogation condition by Shaw et al., it had not been clear whether Nrp items, although recalled significantly better than Rp– items, might not also suffer some ill effects from the retrieval practice given to items in the other categories. In a more recent study of retrieval-induced forgetting effects in an eyewitness context, M. D. MacLeod (2002) also obtained this same pattern of results. Thus, it seems clear that unpracticed items of unpracticed categories escape the potential negative consequences of the retrieval practice given to items belonging to different categories. Additionally and importantly, the results from these two studies demonstrate that retrieval-induced forgetting (i.e., the significant impairment in recall of Rp– items relative to Nrp items)

cannot be attributable to any increase in Nrp recall in the retrieval practice condition.

The results of the Shaw et al. (1995) and M. D. MacLeod (2002) studies have both encouraging and problematical implications for the legal profession. On the one hand, they provide evidence that even minimal interrogation can have a substantial positive impact on the witness's later ability to recall the items that were the targets of that interrogation. On the other hand, they provide evidence that the repeated retrieval of certain details during interrogation might lead to the witness having impaired recall for other details; in particular, details that were not asked about but that happen to bear a category similarity to those items that were the subject of the interrogation. Moreover, the witness's ability to recall such unpracticed details can be significantly impaired compared to that of witnesses who were never interrogated at all. Although, as Shaw et al. cautioned, care should be taken in drawing inferences from these results to the interrogations occurring in actual eyewitness settings, there is no compelling reason to believe that such real-world interrogations would produce a different pattern of results. Indeed, as these investigators argued, the potential consequences – both positive and negative – resulting from the repeated questioning of witnesses in real-life cases could well be greater than those that can be produced in the laboratory.

Although the unintended forgetting by witnesses of unpracticed details as a consequence of the repeated questioning for other details would certainly be a negative consequence of retrieval inhibition, at least the errors so produced would be ones of omission. A potentially more problematical consequence of retrieval inhibition may be that it plays a role in the production of the so-called misinformation effect originally observed by E. F. Loftus and colleagues (e.g., Loftus, 1979; Loftus, Miller, & Burns, 1978; Loftus & Loftus, 1980), which refers to the introduction of errors into the reports of witnesses as a consequence of being presented with misleading information during post-event questioning. The possibility that retrieval-induced forgetting might play such a role was recently explored in research by Saunders and MacLeod (2002), whose reasoning was as follows. In the typical misinformation paradigm, participants are asked questions about, or are given a type of retrieval practice for, a subset of the items or details that they presumably encoded about a previously witnessed event. Consequently, on the basis of the earlier discussed results of Shaw et al. (1995) plus those of MacLeod and Macrae (2001) indicating that retrieval-induced forgetting can be produced by a minimal amount of retrieval practice, such questioning could result in the inhibition of related, but unpracticed, details about the same event. If post-event misinformation about such inhibited details was then introduced, the only information accessible to conscious retrieval or inspection during a later recall attempt would be that misleading post-event material, resulting in the production of misinformation effects.

To test this possible explanation of misinformation effects, Saunders and MacLeod employed an adaptation of the Anderson et al. (1994) retrieval

practice procedure similar to that used by Shaw et al. (1995), but with two additional phases necessary for the introduction of misleading post-event information and the assessment of the misinformation effect. Rather than viewing slides, participants were given two narratives to read about two separate burglaries, one concerning a theft at the Joneses' house, in which 10 items were stolen (e.g., a mobile phone, a Game Boy, a wristwatch) and the other about a theft in the Smiths' house, in which 10 different items were stolen (e.g., a video recorder, rollerblades, a necklace). Following the reading of these narratives, half of the participants were assigned to a retrieval-practice condition in which they were presented with questions about half the items stolen from one of the houses (e.g., the Smiths' house), but were asked no questions about items stolen from the other house, thereby creating Rp+, Rp–, and Nrp items. Similar to the Shaw et al. study, three distinct questions were asked about each of the five practiced items in separate questioning sessions, with the questions increasing in difficulty across the three sessions, and successive sessions separated by distractor tasks. Then, after another distractor task, participants were asked to recall all the items that they could remember about both thefts as a manipulation check to ensure that retrieval-induced forgetting had actually occurred. Following this recall test and another distractor task, participants were asked 12 additional questions about the two burglaries, one of which contained misinformation about one of the stolen items presented in the original narrative (e.g., *necklace* was replaced by *earrings*) and this critical question was directed at either an Rp+, an Rp–, or an Nrp item. Then, following another distractor task, all participants engaged in a final test consisting of multiple-choice questions about the stolen items, with each question presenting three alternatives: the correct item and two new incorrect ones. One of these questions concerned the stolen item about which misleading information had been presented, and for this critical question, the three alternatives were: the original item, the erroneous misinformation, and one new incorrect item. How participants responded to this question was used to determine the occurrence, or lack thereof, of a misinformation effect.

As in the Shaw et al. (1995) study, the other half of the participants did not receive retrieval practice about any of the stolen items, engaging, instead, in retrieval practice for the names of capital cities. They did, however, engage in all other phases of the study, including the misinformation phase and the final multiple-choice test, permitting their performance to serve as an important baseline for the presence of misinformation effects in the absence of related retrieval practice.

Having first determined that retrieval-induced forgetting was produced by the first questioning phase, Saunders and MacLeod (2002) then assessed the consequences of introducing misleading post-event information about items that had been inhibited (i.e., the Rp– items) versus items that had not been inhibited (i.e., Rp+ items, or Nrp items, or control-condition items). If correct in their hypothesis that retrieval-induced forgetting leads to

misinformation effects, such effects should only occur, or at least occur to a greater degree, when the misleading information was presented about inhibited (or Rp–) items versus the other types of non-inhibited items, and, indeed, this pattern of results was exactly what they observed. Significantly more misinformation was reported when such misleading information had been presented about Rp– items compared to all other types of items. Specifically, the likelihood of reporting misinformation was .60 for participants given misinformation about an Rp– item versus .16, .20, and .24 for participants given misleading information about Rp+, Nrp, or control items, respectively.

These results thus provide compelling support for the suggestion of Saunders and MacLeod (2002) that retrieval-induced forgetting could play a critical role in the production of misinformation effects. When participants were asked about or given retrieval practice on a subset of details about a crime event (e.g., some of the items stolen in one of the burglaries), other items known about that event suffered inhibition. Participants then given misinformation about those inhibited items were significantly more likely to report that misleading information, rather than the correct information, on the subsequent multiple-choice test than were participants given misleading information about non-inhibited items. In addition to predicting the pattern of results obtained by Saunders and MacLeod, a retrieval-induced forgetting or inhibitory-based account of misinformation effects would also seem to have certain advantages over other explanations that have been advanced in the literature, such as destructive updating (e.g., Loftus, 1979), retrieval competition (e.g., Bekerian & Bowers, 1983), and source monitoring (e.g., Lindsay & Johnson, 1989a). For example, an inhibitory account is not dependent on the misleading information being introduced subsequent to the encoding of the original material (Lindsay & Johnson, 1989b), and it is compatible with the observation that participants in misinformation studies respond as quickly and as confidently to false memories as they do to genuine ones (Loftus, Donders, Hoffman, & Schooler, 1989).

Stereotyping Most individuals can be thought about in two ways: as a member of some social group, and thus considered to possess some of the attributes typically thought of as characterizing members of that group (i.e., stereotypic attributes), and as an individual, and thus considered to possess attributes that one would not expect to be typical of that group (i.e., counter-stereotypic or individuating attributes). Previously, we have suggested that retrieval-induced forgetting may be an important adaptive mechanism through which we can forget or inhibit memories whose influence on our behaviour we may want to avoid, such as traumatic or disturbing memories or our stereotypic beliefs about others (E. L. Bjork et al., 1998; Macrae & MacLeod, 1999). On the negative side, however, retrieval-induced forgetting may contribute to the well-documented resistance of stereotypes to modification or change (e.g., Hamilton & Sherman, 1994; Johnston & Macrae, 1994)

even in the face of disconfirming evidence (Bodenhausen & Macrae, 1998). If, when we think about members of another group, it is the stereotypic information that more readily comes to mind, then with each such retrieval, the counterstereotopic information will be inhibited, making the stereotype more and more accessible and perpetuating its ability to influence our thinking about such group members. Similarly, retrieval-induced forgetting may be the mechanism by which we maintain other types of pre-existing beliefs despite evidence to the contrary, even beliefs about ourselves.

In recent research, Dunn and Spellman (2003) have investigated some of these possibilities by investigating whether retrieval-induced forgetting could produce stereotype inhibition. For example, in one condition of one of their studies, participants learned about two women, *June* and *Cheryl*, who were presented via the use of pictures as an Asian-American woman and a White mother, respectively, and participants studied six traits associated with each individual, three of which were related to the stereotype about each woman's group (Asian-American traits for June and "mom" traits for Cheryl) and three of which were unrelated to their respective stereotypes. After studying all six traits for each woman, participants were divided into two groups: one given retrieval practice on only the stereotype traits and one given retrieval practice on only the individuating traits. All participants then engaged in a distractor phase, followed by a final test phase in which they were asked to recall all the traits originally studied with each woman. Given these two groups, Dunn and Spellman could ask whether strengthening stereotype traits through retrieval practice would impair access to individuating traits and, conversely, whether strengthening individuating traits would impair access to stereotypic traits. What they found was that retrieval practice on either type of trait produced inhibition or forgetting of the other type of trait.

Perhaps, then, retrieval-induced forgetting provides a more promising way to reduce the influence of stereotypes than seems to be the case for more intentional efforts to do so (e.g., Macrae, Bodenhausen, Milne, & Ford, 1997; Wyer, Sherman, & Stroessner, 2000). In the research by Macrae et al., for example, stereotype-congruent memories were found to be particularly difficult to inhibit intentionally and not to show typical directed-forgetting effects when participants were instructed to forget such information. In contrast, in the Dunn and Spellman study, stereotypic traits were inhibited as an unintended consequence of the retrieval practice given to the individuating traits. On the other hand, as pointed out by Dunn and Spellman, their results do not paint an entirely rosy picture. Although the degree of retrieval-induced forgetting obtained in their study was not significantly moderated by the type of traits given retrieval practice, there was a slight trend suggesting that retrieving stereotype traits inhibited individuating traits more than vice versa. Furthermore, a regression analysis showed that the stronger the participant's belief in the relevant stereotype, the less the stereotype traits were inhibited by practice of the individuating traits. As suggested by Dunn

and Spellman, one possible explanation for this relationship could be that stereotype traits comprise a more highly integrated set for strong believers, and, thus, may be able to escape the effects of inhibition via mediated retrieval. That is, being strongly linked together, if one trait manages to be retrieved, it can provide a retrieval path to the others (see the earlier discussion of Anderson & McCulloch, 1999).

Additionally, as the results of M. D. MacLeod and Macrae (2001) have shown, retrieval-induced forgetting may be fairly transient, lasting for at least 20 minutes but perhaps no longer than 24 hours. Typically, given our constantly changing needs and goals, the temporary nature of such forgetting is advantageous. The information we inhibit today may be exactly what we need tomorrow, next week, or next year. With respect to some stereotypes, however, a more permanent state of inhibition may be desirable. But to keep such information in a more or less permanent state of inhibition may well require recurring practice of the counterstereotypic or individuating traits.

The role played by retrieval-induced forgetting *vis-à-vis* stereotypes is thus complex. If we mainly practice retrieving stereotypic congruent information, our stereotypes will be strengthened while our access to counterstereotypic or individuating information will become more and more inhibited – clearly, an unintended negative consequence of retrieval-induced forgetting. On the other hand, if we consciously practice retrieving the individuating information, we will automatically inhibit the stereotypic information, making it less and less likely to come to mind – clearly, an unintended positive consequence of retrieval-induced forgetting.

Suppression of negative memories Another area in which the unintended consequences of retrieval-induced forgetting could potentially serve a positive function is in the suppression of various types of unpleasant or negative memories, thereby enabling us to avoid the potentially disabling emotions that would be aroused by the continuous presence in consciousness of such memories. For example, as we have previously suggested (E. L. Bjork et al., 1998; Macrae & MacLeod, 1999), in many abusive situations, it seems reasonable to assume that the victim of the abuse would hold both positive and negative memories about the perpetrator of the abuse. It also seems reasonable to assume that, in most such situations, there would be reasons for the victim of abuse to want to retrieve only the positive memories associated with the perpetrator, rather than dwelling on the negative ones. If so, then when thinking about the perpetrator, the victim would engage in retrieval practice for the positive memories. Thus, those memories – like the practiced exemplars of the Anderson et al. (1994) paradigm – would become more and more likely to be retrieved in the future, while the negative memories – being selected against and thus inhibited again and again – would become less and less accessible to retrieval. And, as long as the practice of the positive continued, access to the negative would remain inhibited.

Although more research is necessary to demonstrate such a consequence for retrieval-induced forgetting, some initial results point in that direction. We have found, for example, that having participants practice neutral traits associated with fictional individuals seems to be making their valenced traits (either positive or negative) less accessible for retrieval, and, furthermore, that the traits that are otherwise most recallable are the most susceptible to such effects (Storm, E. L. Bjork, & R. A. Bjork, 2005). Although few of us may have been actual victims of serious abuse, essentially all of us must interact with others whom we consider to have both positive and negative characteristics, or with whom we have had both positive and negative interactions. Thus, being able to remember or emphasize such people's positive attributes while forgetting or making their negative attributes less salient in our memories, can help us to maintain positive relationships with such people – who can be our co-workers, our supervisors, members of our family, or even our spouses – that is, people with whom we want or need to maintain functional and positive relationships. Storm et al.'s results suggest that these types of unintended, but nonetheless positive, consequences of retrieval-induced forgetting could well serve as an adaptive mechanism to smooth our interactions with people with whom we must interact on a range of levels and in a variety of contexts.

CONCLUDING COMMENTS

We began this chapter with the assertion that forgetting, contrary to our intuitions, is a necessary and critical component of an efficient and adaptive memory system. More specifically, we have focused on two types of forgetting – intentional (or "directed") forgetting and retrieval-induced forgetting – each of which, in our view, plays an important role in the everyday functioning of our memories. It is important to emphasize, however, that the benefits of each type of forgetting, with respect to keeping our memories current and avoiding competition from competing memories, are accompanied by some potential costs. Each type of forgetting also has a dark side, so to speak – beyond simply making information unavailable that we might then later, for whatever reason, want to recall. In the case of intentional forgetting, our opinions and beliefs may continue to be influenced by intentionally forgotten information in ways about which we are unaware and thus cannot mitigate. In the case of retrieval-induced forgetting, for example, we may unintentionally be led to accept misleading information as correct or to maintain certain stereotypical beliefs that we would prefer not to maintain. Forgetting, in short, refines and updates our memories, but not without certain costs.

References

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

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

Anderson, M. C., Green, C., & McCulloch, K. C. (2000). Similarity and inhibition in long-term memory: Evidence for a two-factor theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 1141–1159.

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

Anderson, M. C., & Neely, J. H. (1996). Interference and inhibition in memory retrieval. In E. L. Bjork & R. A. Bjork (Eds), *Handbook of perception and cognition: Vol. 10. Memory* (pp. 237–313). San Diego: 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.

Basden, B. H., Basden, D. R., & Gargano, G. J. (1993). Directed forgetting in implicit and explicit memory tests: A comparison of methods. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 603–616.

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

Bjork, E. L., & Bjork, R. A. (1996). Continuing influences of to-be-forgotten information. *Consciousness and Cognition*, *5*, 176–196.

Bjork, E. L., & Bjork, R. A. (2003). Intentional forgetting can increase, not decrease, residual influences of to-be-forgotten information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 524–531.

Bjork, E. L., Bjork, R. A., & Anderson, M. C. (1998). Varieties of goal-directed forgetting. In J. M. Golding & C. N. MacLeod (Eds.), *Intentional forgetting* (pp. 103–137). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Bjork, E. L., Bjork, R. A., & Glenberg, A. M. (1973, November). *The reinstatement of proactive interference owing to to-be-forgotten items*. Paper presented at the meeting of the Psychonomic Society, St Louis, MO.

Bjork, E. L., Bjork, R. A., & Kilpatrick, H. A. (1990, November). *Direct and indirect measures of inhibition in directed forgetting*. Paper presented at the meeting of the Psychonomic Society, New Orleans, LA.

Bjork, E. L., Bjork, R. A., Stallings, L., & Kimball, D. R. (1996, November). *Enhanced false fame owing to instructions to forget*. Paper presented at the meeting of the Psychonomic Society, Chicago, IL.

Bjork, R. A. (1972). Theoretical implications of directed forgetting. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory*. Washington, DC: Winston & Sons.

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

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

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

Block, R. A. (1971). Effects of instructions to forget in short-term memory. *Journal of Experimental Psychology, 89*, 1–9.

Bodenhausen, G. V., & Macrae, C. N. (1998). Stereotype activation and inhibition. In R. S. Wyer, Jr. (Ed.), *Sterotype activation and inhibition. Advances in social cognition* (Vol. 11, pp. 1–52). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

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

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

Elmes, D. G., Adams, C., & Roediger, H. L. (1970). Cued forgetting in short-term memory: Response selection. *Journal of Experimental Psychology, 86*, 103–107.

Geiselman, R. E., & Bagheri, B. (1985). Repetition effects in directed forgetting: Evidence for retrieval inhibition. *Memory & Cognition, 13*, 51–62.

Geiselman, R. E., Bjork, R. A., & Fishman, D. (1983). Disrupted retrieval in directed forgetting: A link with posthypnotic amnesia. *Journal of Experimental Psychology: General, 112*, 58–72.

Hamilton, D. L., & Sherman, J. W. (1994). Stereotypes. In R. S. Wyer, Jr., & T. K. Srull (Eds.), *Handbook of social cognition: Vol. 2. Applications* (pp. 1–68). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Johnson, H. (1994). Processes of successful intentional forgetting. *Psychological Bulletin, 116*, 274–292.

Johnston, L. C., & Macrae, C. N. (1994). Changing social stereotypes: The case of the information seeker. *European Journal of Social Psychology, 24*, 581–592.

Kihlstrom, J. F., & Cantor, N. (1984). Mental representations of the self. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 17, pp. 1–47). New York: Academic Press.

Klein, S. B., & Kihlstrom, J. F. (1986). Elaboration, organization, and the self-reference effect in memory. *Journal of Experimental Psychology: General, 115*, 26–38.

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

Lindsay, D. A., & Johnson, M. (1989a) The eyewitness suggestibility effect and memory for source. *Memory & Cognition, 17*, 349–358.

Lindsay, D. A., & Johnson, M. (1989b) The reversed suggestibility effect. *Bulletin of the Psychonomic Society, 27*, 111–113.

Loftus, E. F. (1979). *Eyewitness testimony*. Cambridge, MA: Harvard University Press.

Loftus, E. F., Donders, K., Hoffman, H. G., & Schooler, J. W. (1989). Creating new memories that are quickly accessed and confidently held. *Memory & Cognition, 5*, 607–616.

Loftus, E. F., Miller, D. G., & Burns, H. J. (1978). Semantic integration of verbal information into a visual memory. *Journal of Experimental Psychology: Human Learning and Memory, 4*, 19–31.

Loftus, G. R., & Loftus, E. F. (1980). On the permanence of stored information in the human brain. *American Psychologist, 35*, 409–420.

MacLeod, C. M. (1998). Directed forgetting. In J. M. Golding & C. M. MacLeod

(Eds.), *Intentional forgetting: Interdisciplinary approaches* (pp.1–57). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

MacLeod, M. D. (2002). Retrieval-induced forgetting in eyewitness memory: Forgetting as a consequence of remembering. *Applied Cognitive Psychology, 16*, 135–149.

MacLeod, M. D., Bjork, E. L., & Bjork, R. A. (2003). The role of retrieval-induced forgetting in the construction and distortion of memories. In B. Kokinov & W. Hirst (Eds.), *Constructive memory: NBU Series in cognitive science* (pp. 55–68). Sophia: New Bulgarian University.

MacLeod, M. D., & Macrae, C. N. (2001). Gone but not forgotten: The transient nature of retrieval-induced forgetting. *Psychological Science, 121*, 148–152.

Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Ford, R. L. (1997). On the regulation of recollection: The intentional forgetting of stereotypical memories. *Journal of Personality and Social Psychology, 72*, 709–719.

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

Macrae, C. N., & Roseveare, T. A. (2002). I was always on my mind: The self and temporary forgetting. *Psychonomic Bulletin & Review, 9*, 611–614.

McGeoch, J. A. (1942). *The psychology of human memory*. New York: Longman.

Mensink, G. J. M., & Raajimakers, J. W. (1988). A model of interference and forgetting. *Psychological Review, 95*, 434–435.

Moeser, S. D. (1979a). Retrieval interference in the independent storage of related episodic traces. *Canadian Journal of Psychology, 33*, 182–192.

Moeser, S. D. (1979b). The role of experimental design in investigations of the fan effect. *Journal of Experimental Psychology: Human Learning and Memory, 5*, 125–134.

Postman, L. (1971). Transfer, interference and forgetting. In J. W. Kling & L. A. Riggs (Eds.), *Woodworth and Scholsberg's experimental psychology* (3rd ed., pp. 1019–1132). New York: Holt, Rinehart & Winston.

Raajimakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review, 88*, 93–134.

Reed, H. (1971). Studies of the interference processes in short-term memory. *Journal of Experimental Psychology, 84*, 452–457.

Roediger, H. L., (1974). Inhibiting effects of recall. *Memory & Cognition, 2*, 261–269.

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

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

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

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

Smith, R. E., & Hunt, R. R. (2000). The influence of distinctive processing on retrieval-induced forgetting. *Memory & Cognition, 28*, 503–508.

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

Wells, G. L. (1993). What do we know about eyewitness identification? *American Psychologist, 48*, 553–571.

Wyer, N. A., Sherman, J. W., & Stroessner, S. J. (2000). The roles of motivation and ability in controlling the consequences of stereotype suppression. *Personality & Social Psychology Bulletin*, *26*, 13–25.