



Rethinking interference theory: Executive control and the mechanisms of forgetting

Michael C. Anderson*

Department of Psychology, 1227 University of Oregon, Eugene, OR 97403-1227, USA

Received 21 April 2003; revision received 13 August 2003

Abstract

Interference provides an account of one of the most basic problems in the science of memory: forgetting. Historically, theories of this process were shaped by models of associative learning prevalent when interference research began. In this article, I argue that we should reconsider the long-standing conceptualization of interference as a learning phenomenon and reframe interference as arising from systems that achieve mental and behavioral control. Specifically, it is argued that forgetting is not a passive side effect of storing new memories, but results from inhibitory control mechanisms recruited to override prepotent responses. In support of this idea, I discuss two control situations in which response override is necessary—selection and stopping—and show how these situations have direct parallels in retrieval. I then review evidence that in both of these situations, the need to override prepotent, distracting memories is supported by inhibitory mechanisms that ultimately cause forgetting. The theoretical properties of these inhibitory effects are outlined, along with critical factors known to modulate or mask inhibition. The relation between this executive control theory of forgetting and classical accounts of interference is discussed.

© 2003 Published by Elsevier Inc.

Over a century ago, G. E. Mueller and Pilzecker reported one of the first empirical demonstrations of forgetting due to interference (Müller & Pilzecker, 1900). In this classic work, Mueller and Pilzecker found that people were less likely to recall a memory item if in the interim the retrieval cue that was used to test that item had become associated to another memory. They named this effect retroactive inhibition, highlighting the manner in which the storage of new experiences interferes with memories encoded earlier in time. Mueller and Pilzecker believed that this memory impairment occurred because the process of storing new memories disrupted the consolidation process that would have ordinarily strengthened the traces that subjects had acquired earlier. By this view, all would-be memories perseverate for a brief period after they are encoded, as evidenced by the

tendency of recent memories to pop into consciousness unbidden by any particular cue. Perseveration was thought to be necessary to more firmly fix a trace into long-term storage. If another effortful activity intervened (such as learning a second list of items), the perseverative process for the earlier memories was thought to be dampened, ultimately preventing the traces from being woven into the fabric of memory.

Although the disrupted consolidation theory was largely abandoned as an account of retroactive interference (see McGeoch & Irion, 1952, for arguments), the phenomenon itself and the method Mueller and Pilzecker introduced to study it have played a central role in shaping the history of memory research. Their work set off the classical interference era (1900–1970) in memory research. In this era, considerable energy was devoted to unraveling the mechanisms of interference—a focus deemed worthy because it addressed the fundamental problem of forgetting. How is it possible for an experience that is vivid and lively in our memories today to

* Fax: 1-541-346-4911.

E-mail address: mcanders@darkwing.uoregon.edu.

ultimately fade? Why do memories grow less accessible over time? According to classical interference theory, Mueller and Pilzecker's discovery provided an answer: it was not the passage of time that caused forgetting, as might seem intuitive, but rather, changes correlated with time, such as the storage of new experiences into memory—in particular, highly similar experiences—that were at the root of memory failures. You can remember what you had for dinner yesterday evening now, yet in a few months you will not be able to, not because time has passed, but because the routine nature of our lives ultimately clutters memory with many highly similar dinner events. This clutter makes any particular memory very difficult to retrieve. Thus, when we forget, it is not because memories decay, but because we are victims of the ever-changing structure of our memory and of basic limitations in our ability to differentiate similar traces. This view has stood the test of time: after 70 years of research and after tens of thousands of papers on the topic, there can be little doubt that interference is a powerful cause of forgetting.

What can be doubted, however, is the manner in which interference causes forgetting. On one hand, forgetting may be a direct consequence of adding new traces into memory. Both classical and modern theories have emphasized this approach. For instance, McGeoch's influential response competition theory (McGeoch, 1942) attributed interference effects to heightened competition arising from the association of additional traces to a retrieval cue (or to the strengthening of an existing competitor); in his framework, forgetting was a consequence of adding new associative structure. Modern theories such as those embodied in relative strength or ratio-rule models of retrieval (Anderson, 1983; Mensink & Raaijmakers, 1988) are the conceptual descendants of this view in their emphasis on how retrieval of a given item is impeded by competing associations (see Anderson & Bjork, 1994, for a review). Structure-based theories such as these do not require special mechanisms of forgetting and have the virtue of parsimony. On the other hand, they de-emphasize a basic problem in how we use our memory: how do we overcome interference between competing traces to retrieve the memory we want? What are the repercussions, if any, of resolving competition for the traces that interfered? Given that our cognitive goals often require the recall of specific events in long-term memory, some process must exist for resolving interference.

In this article, I present a view of how interference leads to forgetting which emphasizes how interference gets resolved. I argue that a theory of interference should be framed in the larger context of how organisms control the direction of their actions and thoughts. By this view, memory retrieval presents a special case of a broad class of situations that recruit executive control processes; it is the executive control mechanism that overcomes inter-

ference—inhibition—that causes us to forget, not the competition itself. This view departs from the common assumption that forgetting is a passive side effect of the ever-changing structure of memory. The mere storage of interfering traces is not what causes memories to grow less accessible with time. Rather, forgetting, whether incidental or intentional, is produced as a response to interference caused by activated competitors in memory. I review the evidence supporting this functional view. This review focuses on retrieval-induced forgetting and more recent work with the think/no-think paradigm and provides a specification of the theoretical properties of those phenomena, some of their boundary conditions, and empirical challenges to measuring inhibition. In the final section, I contrast the proposed view with classical theories of interference. Before beginning, however, I elaborate on the theoretical perspective advanced here.

Executive control and the mechanisms of retrieval

The current perspective begins with a simple observation about human behavior: Actions, once started, can usually be stopped. This simple fact was impressed upon me one evening while opening the kitchen window. As the window slid along its track, it pushed a small cactus off the edge of the sill. My hand darted reflexively to catch the falling cactus. Mere centimeters from it, I stopped my hand from clutching the cactus's needle-dense body. The plant dropped to the floor and was ruined, but I was happy to have avoided piercing my hand with thousands of little red needles. This last minute save was made possible by my ability to terminate physical action—an ability so pervasive that it goes nearly unnoticed in daily life.

The preceding case is a classic example of a situation in which we need to overcome a strong habitual response—a situation widely regarded as requiring executive control. This is sometimes referred to as response override, and is illustrated in Fig. 1. In response override, one must stop a prepotent response to a stimulus (such as a falling object). This may either be because the circumstance requires that the response be withheld, or because a less common response is more contextually appropriate. For example, it is more contextually appropriate to say "Hola" when someone waves to you while you are in Spain, even though your habitual response may be to say "Hello." The capacity to either stop or redirect action in this way is crucial to daily life. Without it, we would lose essential flexibility to adapt behavior according to changes in our goals, or to changes in the environment itself. We would be slaves to habit or reflex.

A key theoretical question that this problem raises is "How do we keep from being automatically controlled

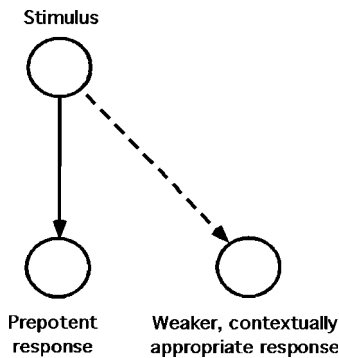


Fig. 1. Graphical depiction of a typical response override situation. Circles correspond to representations in memory, with lines representing associations between these representations. The stimulus is associated to two responses, one of which is stronger (prepotent), and the other of which is weaker (depicted by a dotted line). Response override must occur when either the organism needs to: (1) emit the weaker response, when it is more contextually appropriate, despite the stronger association to the prepotent response, or (2) stop any response from occurring. Inhibitory control is thought to suppress activation of the prepotent response to permit response override. The response override situation characterizes many paradigms in work on executive function, including the Stroop and go/no-go tasks.

by the habitual action?” One widely discussed answer to this question is that response override is accomplished by inhibiting the undesired action. According to this view, the presentation of a stimulus activates a representation corresponding to that stimulus in long-term memory. Activation then spreads from that representation to associated responses in proportion to how strongly associated they are to the stimulus. When a response becomes sufficiently activated, it will be emitted. If a stimulus is associated to multiple responses, the one that achieves threshold most quickly will generally be emitted, pre-empting other responses. However, if a weaker response is more contextually appropriate, inhibition can be recruited to suppress the stronger one. Inhibition is thought to reduce the level of activation for a given response, preventing it from achieving threshold. In so doing, this process permits weaker, but more contextually appropriate responses to be expressed, enabling flexible, context-sensitive behavior. This is known as inhibitory control.

Given the putative importance of inhibitory control in directing overt behavior, it is reasonable to ask whether internal actions might also be the target of such mechanisms. Clear parallels exist between the control of action and the control of memory. Just as a stimulus may spread activation to a prepotent motor response, a retrieval cue may spread activation to a strongly associated item in memory, leading it to be retrieved.

The retrieval of associated memories is not always desirable; sometimes, we may wish to retrieve another memory item that is associated to the cue driving retrieval, but that is more weakly associated to that cue; other times, we may wish to avoid retrieval altogether either because the associated memory is unpleasant or simply because we wish to maintain conceptual focus on the concept that is acting as a cue. Although we sometimes retrieve things that we do not intend, we often are able to exercise control over this tendency; we can recollect the particular event we are seeking despite interference from stronger competitors, and we can stop ourselves from thinking about unwanted memories. Given these functional parallels between motor behavior and memory retrieval, it is possible that response override mechanisms are recruited to control unwanted memories that intrude by virtue of spreading activation (for related arguments, see Shimamura, 1995). If so, we should find evidence for inhibitory control in memory situations likely to involve response override—situations such as the need to select a weaker, yet more contextually appropriate response, given interference from one or more prepotent competitors, or the need to stop a response altogether.

A core claim of this article is that strong evidence for these parallels exists, and that inhibitory processes recruited during the control of memory retrieval precipitate the forgetting associated with interference. In support of this executive control perspective, I review evidence for a role of inhibitory processes in memory selection and more briefly, in memory stopping. Memory selection is required during retrieval when our goal is to recall an event or fact from long-term memory in the face of interference from related traces that become activated by cues guiding retrieval. The need to stop retrieval arises when we confront a cue or reminder and we wish to prevent an associated memory from entering awareness. In both situations, attempts to limit the influence of activated and distracting memories have been found to impair their later accessibility, highlighting an important link between forgetting and the control of retrieval. In both cases, the memory impairment is better explained by inhibition than by conventional associative interference mechanisms. The forgetting induced by inhibition is often adaptive, limiting the tendency for outdated or intrusive memories to disrupt performance (Bjork, 1989; see also Anderson, 2001; Anderson & Green, 2001).

Inhibitory control in selective memory retrieval

The need to select a weaker response to a stimulus in the face of interference from a prepotent competitor finds a natural parallel in memory in the situation of selective retrieval. Here, the aim is to recall a particular target event or fact when provided with one or more

retrieval cues. Typically, a given retrieval cue will be associated with other memories as well—and some may be more strongly associated to the cue than the target item. Since the classical interference era, we have known that when multiple traces are associated to the same cue, they tend to compete for access to conscious awareness (see Postman, 1971; see Anderson & Neely, 1996 for reviews). This form of competition presents a problem of control because the retrieval cue by itself cannot be relied upon to access the target item—in fact, the presence of a strong competitor could in principle perpetually divert us from that target memory. If inhibitory control mechanisms are recruited to override prepotent responses, it seems reasonable that they might also be used to override prepotent memories. To the extent that the effects of inhibitory control persist, then situations demanding the selective retrieval of a target item should cause long-lasting memory impairment for suppressed competitors. Thus, the very act of remembering should cause forgetting of related memories.

Over the last decade, we have explored the foregoing prediction with a paradigm we developed to examine the effects of retrieval on related memories: the retrieval practice paradigm (Anderson, Bjork, & Bjork, 1994). In the typical experiment, subjects study lists of category–exemplar pairs (e.g., fruit–orange, drinks–scotch, fruit–banana). They then perform retrieval practice on half of the exemplars from half of the categories by completing cued stem recall tests (e.g., fruit-or_____). Each practiced item is tested three times during the retrieval practice phase to increase the magnitude of the effect on related items. After a 20-min retention interval, subjects are given a final cued recall test for all the exemplars. Performance on this test can be measured for three item types: practiced items (e.g., orange), unpracticed items from the practiced categories (e.g., banana), and unpracticed baseline items from unpracticed categories (e.g., scotch). Fig. 2A illustrates our initial findings with this paradigm, which are quite typical. As can be seen, recall of the practiced exemplars was improved on the final test relative to performance on baseline categories, demonstrating the well documented benefits of retrieval-practice on the practiced items themselves (Allen, Mahler, & Estes, 1969; Bjork, 1975; Carrier & Pashler, 1992; Gardiner, Craik, & Bleasdale, 1973).

However, recall for the unpracticed exemplars from the practiced categories (e.g., banana) was significantly worse than for the items from baseline categories (e.g., drinks). Thus, remembering some items during the retrieval practice phase caused subjects to forget other things that were related to them on a delayed retention test 20 min later. We have referred to this finding as *retrieval induced forgetting* (Anderson et al., 1994), to highlight the central role that retrieval is believed to play in generating the effect. Research on retrieval-induced forgetting builds on classic work on the phenomenon of

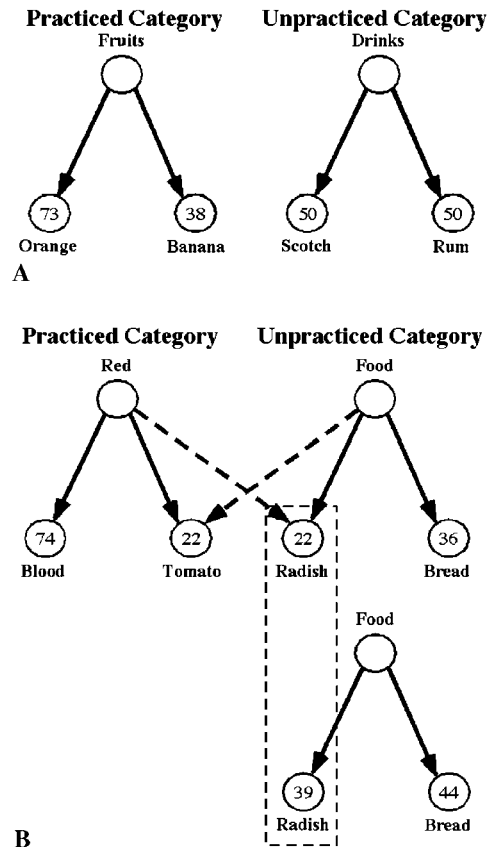


Fig. 2. Stimulus structures in two typical retrieval-induced forgetting experiments with representative results. (A) A typical within-category retrieval-induced forgetting study, as done by Anderson et al. (1994). The example illustrates two items from each of two categories that subjects have studied (six items are usually studied in eight categories), for purposes of illustration. In this example, subjects have performed retrieval practice on Fruits Orange, but not on Fruits Banana (unpracticed competitor) or any members from the Drinks category (an unpracticed baseline category). As shown here, practice typically facilitates recall of the practiced item, and impairs recall of the unpracticed competitor, relative to performance in baseline categories. (B) Stimulus structure and results from a typical cross category inhibition experiment, as performed by Anderson and Spellman (1995). In the related condition (top half of (B)), subjects study two related categories (Red Things and Foods) and then perform retrieval practice on some of the members of one of them (e.g., Red Blood), but not the other (Foods). As shown in (B), this not only impairs the delayed recall of unpracticed competitors that are explicitly studied under the Red category (e.g., Red-tomato), but also those competitors (i.e., other Red things) that are studied and tested under a separate category (e.g., Food Radish). This can be seen by comparing performance to items in the corresponding category (i.e., Food) when the red category is not studied or practiced (i.e., the Unrelated condition; see dotted box for the appropriate comparison). The impairment of items in a separate category is an example of the cue-independence of inhibition.

output interference (Arbuckle, 1966; Dong, 1972; Roediger, 1973; Roediger & Schmidt, 1980; Smith, 1971, 1973; Smith, D'Agostino, & Reid, 1970; Tulving & Arbuckle, 1963, 1966) in which it was shown that the probability of recalling a previously studied item declines with the serial position of that item in the testing sequence. However, work on retrieval-induced forgetting establishes that retrieval-related memory impairments can be long-lasting, and are not limited to dynamics occurring in a single "output" session. Crucially, retrieval-induced forgetting is consistent with the view that inhibitory control mechanisms are recruited to overcome interference during retrieval practice, with inhibition manifesting as recall impairment for competitors on the final retention test.

Although retrieval-induced forgetting could be produced by inhibition, the preceding results do not clearly establish inhibition as the mechanism. The basic finding of retrieval-induced forgetting is compatible with McGeoch's (1942) classical response competition theory of interference. According to McGeoch's theory, the likelihood of recalling a target response should go down either when a new response gets associated to the cue normally used to retrieve it, or when an existing alternative response is strengthened. In either case, the target item will suffer increased competition from the alternative response—competition that will block access to that target. These competitive dynamics have become formalized in several modern memory architectures that posit relative strength theories of retrieval (e.g., Anderson, 1983; Raaijmakers & Shiffrin, 1981). In these models, the probability of recalling a target is determined by that item's strength of association to a cue, relative to the strengths of association of all items related to that cue. Thus, when an alternative response is strengthened, say by retrieval practice, the relative strength of all other nonpracticed items declines on subsequent tests. Later, when the subject tries to recall the target, the strengthened competitor will have a retrieval advantage that will lead it to intrude so persistently that subjects will abandon their efforts to recall the unpracticed exemplars (see also, Rundus, 1973). Importantly, this approach does not require inhibition; rather, practiced items become so strongly linked to the practice cue that they block the retrieval of other exemplars. This blocking account is plausible, given the substantial strengthening that practiced items typically enjoy (however see later section on strength independence). Other noninhibitory mechanisms may also contribute to retrieval induced forgetting. For example, retrieval practice may damage the association linking the category to the affected exemplar or alter instead the meaning of the practiced category cue (e.g., by biasing "Fruits" towards "Citrus fruits") so that the category label is no longer a functional cue for retrieving the unpracticed competitor. All of these mechanisms have

been proposed as theories of interference (for a review of non-inhibitory sources of memory impairment, see Anderson & Bjork, 1994; see description of a subset of these in the later discussion). Although it might seem difficult to distinguish these alternatives, focused empirical research has yielded evidence favoring the inhibitory control view. I discuss this next.

Properties of retrieval-induced forgetting

Work on retrieval-induced forgetting has revealed properties that uniquely support the inhibitory control hypothesis, and that suggest that alternative strength-based models may not be correct. These include cue-independence, retrieval-specificity, interference-dependence, and strength-independence. I discuss these properties next, along with other findings that are of theoretical interest.

Cue-independence

Many theories of interference predict that forgetting should be strongly cue-dependent—that is, observations of forgetting should be tied to a particular cue. For example, the blocking theory asserts that strengthening some exemplars through retrieval practice (e.g., Fruit Orange) impairs the recall of other exemplars (e.g., Fruit Banana) on a delayed test because the presentation of their shared cue at test leads the stronger response (orange) to intrude persistently and block the weaker item. If, however, one were to try to recall the weaker item through an independent test cue not associated to the practiced item (e.g., Monkey B___), associative competition should be circumvented. Thus, whether one observes forgetting of Banana should depend on whether one uses the original retrieval practice cue to test the critical item or an independent cue. Theories that propose that interference derives from unlearning of the cue-target association between the practiced category and the critical item, diversion of activation resources, or biases in the meaning of the retrieval practice cue all share this same feature of predicting cue-dependent forgetting (see Anderson & Bjork, 1994; Anderson & Spellman, 1995, for discussion).

The inhibitory control perspective, by contrast, predicts that retrieval induced forgetting should exhibit cue-independence—that is, a tendency for the impairment to generalize to novel test cues not involved in the retrieval practice events that caused impairment. This prediction follows because impairment is thought to arise from suppression of the competing memory itself, rather than from damage to any particular association. Thus, performing retrieval practice on Fruit-Orange should reduce activation for the item Banana. If Banana is less active, it should not matter whether the item is tested from the original retrieval practice cue (Fruit) or from a novel test cue (e.g., Monkey B___). To test this, Anderson

and Spellman (1995) modified the retrieval practice paradigm for use with new materials (Fig. 2B). As in the original retrieval-induced forgetting experiment, subjects studied categories containing six exemplars each, but unlike in that study, the categories were related. For example, although subjects studied Blood and Tomato under the Red category, Tomato is also a Food; and although Radish and Crackers were studied as Foods, Radish is also a Red thing. The key question was whether retrieval practice on items such as Red-Blood would not only impair competitors explicitly studied under the same category cue, like Red-Tomato, but also red things like Radish that were studied and tested under a separate category cue.

According to the response competition view, retrieval practice on Red-Blood should not impair delayed recall for Food Radish, even if retrieval practice strengthens the Red-Blood association and weakens the Red-Radish association. Radish should remain unimpaired because it is tested with the Food category—a different retrieval cue that circumvents those factors. However, if retrieval practice on Red-Blood initially activates all of the Red items, both Tomato and Radish should become activated, causing interference that triggers inhibitory control. The resulting suppression of Radish should be observable later when it is tested with Food. As Fig. 2B shows, the recall of Food–Radish was impaired.

These data show that inhibitory processes contribute to retrieval-induced forgetting, rendering competing memories less accessible regardless of which cue is used to test them. Evidence for cue-independent forgetting has now been found many times with stimuli varying in both type and complexity (Anderson & Bell, 2001; Anderson & Green, 2001; Anderson, Green, & McCulloch, 2000; Anderson et al., submitted; Anderson & Shivde, in preparation a; Anderson & Shivde, in preparation b; Anderson & Spellman, 1995; Johnson & Anderson, in press; Levy, Reinholz, & Anderson, in preparation; Miyamoto & Anderson, in preparation; Radvansky, 1999; Shivde & Anderson, 2001; however, see Williams & Zacks, 2001). Taken as a whole, these findings show that cue-independence is a general property of retrieval-induced forgetting and that cue dependent mechanisms such as blocking are not adequate to account for the effect.

Retrieval specificity

Although cue-independence argues that inhibition causes retrieval-induced forgetting, cue-dependent forgetting mechanisms may nevertheless contribute in some cases. In particular, whenever the retrieval practice cue is used during later tests of subjects' memory, both inhibitory and noninhibitory mechanisms may cause forgetting. For instance, the impairment of Red-Tomato in the preceding example (an item that was both studied and tested with the retrieval practice cue) may reflect a

mixture of blocking and inhibition effects. As can be seen in Fig. 2B, however, this seems unlikely. If two sources of forgetting contributed to impairment on items such as Red-Tomato, but only one source (inhibition) on Food-Radish, we would expect to see more retrieval-induced forgetting in the former instance than in the latter. The failure to find such a difference casts doubt on the role of blocking in retrieval-induced forgetting even when the practiced category is used as a test cue. Nevertheless, a more direct test of the role of blocking in retrieval-induced forgetting would be desirable.

According to the blocking hypothesis, presenting the retrieval practice category on the final test leads practiced items to intrude perseveratively, blocking recall of the unpracticed competitors. If so, then strengthening practiced items in any way should impair related competitors. Impairment should be found, for example, even if items are strengthened with repeated study exposures instead of retrieval practice. Several studies have addressed this possibility. For example, using Anderson and Spellman's cross-category inhibition paradigm, Anderson and Shivde (in preparation a) manipulated whether the to-be-practiced items were strengthened by retrieval practice or repeated study exposures. The retrieval-practice condition replicated both the within and cross-category impairment observed by Anderson and Spellman (1995). Extra study exposures, however, failed to impair related items. No inhibition was found despite the fact that both strengthening methods facilitated the practiced items to the same degree, as evidenced by the substantial increase in their recall on the final test. Several investigators have found this pattern, using a variety of different types of materials and different paradigms (Anderson & Bell, 2001; Anderson, Bjork, & Bjork, 2000; Bäuml, 1996, 1997, 2002; Blaxton & Neely, 1983; Ciranni & Shimamura, 1999; Shivde & Anderson, 2001). Taken together, these findings argue that associative blocking does not contribute substantially to within-category retrieval-induced forgetting, nor perhaps to interference effects more broadly. Rather, inhibition is driven by the need to override interference from competing memories during the selective retrieval of target items.

Interference dependence

Retrieval may be necessary to induce inhibition, but it is not sufficient. According to the executive control theory, retrieval induced forgetting should only arise whenever a related memory interferes with the retrieval of a target item and triggers inhibitory control. If a related item does not interfere, it should not be inhibited even when a target has been retrieved.

Several studies favor the view that retrieval-induced forgetting is moderated by the amount of interference caused by a competing item. For instance, Anderson et al. (1994) found that retrieval practice did not always

impair the later recall of related exemplars. Retrieval practice primarily caused impairment when related category exemplars were high in taxonomic frequency (e.g., Fruit Banana). Low frequency competitors (e.g., Fruit Guava) were always less impaired and often exhibited no measurable impairment at all, even when subjects performed retrieval practice on exactly the same items. Anderson et al. (1994) found that the main determinant of the amount of retrieval-induced forgetting was neither the taxonomic frequency of the practiced items, nor the degree to which practiced items were strengthened on the final test, but rather the frequency of the competitors. The more strongly associated to the category an unpracticed competitor was, the more impairment was found. The same pattern has been found in an output interference design: high, but not low taxonomic frequency exemplars exhibit within category output interference (Bäuml, 1998; see also, Anderson et al., 1994, Experiment 2; however, see Anderson et al., 1994, Experiment 3). These findings are consistent with the idea that inhibitory control is most needed when a related item interferes during retrieval, as might be expected of the most dominant exemplars.

Interference dependence has been demonstrated in other ways, as well. For example, retrieval induced forgetting can be eliminated simply by manipulating the interference demands of the retrieval practice task. This was demonstrated by Anderson et al. (2000). In their competitive retrieval practice group, subjects were given the category and the first two letters of an exemplar as cues (e.g., Fruit Or__ for Orange) during each practice trial (as is typically done), and subjects were asked to recall the item they had studied. In the non-competitive practice condition, subjects also performed retrieval practice, but on the category name. Specifically, subjects were given the first two letters of the category name, with an exemplar (e.g., Fr__ Orange for Fruit Orange), and were asked to recall the category name. Anderson et al. (2000) argued that related exemplars were unlikely to interfere with the retrieval of the category name because a practiced exemplar itself, which was associated to the category and not with the other exemplars, served as a retrieval cue. As predicted, Anderson et al. (2000) found inhibition in the competitive, but not in the noncompetitive condition. This difference was found despite the presence of retrieval in both conditions and despite significant and comparable strengthening of practiced items. Thus, when the retrieval task itself does not require interference to be resolved, little retrieval-induced forgetting is found, even when the nature of the competitor is held constant.

In the foregoing studies, the degree to which practiced items were strengthened was nearly identical in both the conditions that showed and did not show retrieval-induced forgetting. These results suggest that inadequate strengthening of practiced items is unlikely

to be the cause of differential impairment. However, one might still be concerned that the degree of strengthening was not manipulated strongly enough to reveal impairment. To address this, Shivde and Anderson (2001) performed a parametric manipulation of the number of retrieval practice trials given to a practiced item, to see whether impairment might emerge for weaker competitors. To manipulate the degree of interference, Anderson and Shivde used asymmetric homographs, pairing each one with one word related to its dominant sense (e.g., Arm Shoulder) and another related to its subordinate sense (e.g., Arm Missile). Subjects were then asked to perform retrieval practice either 0, 1, 5, or 20 times on either the dominant or the subordinate word associate. Following retrieval practice, subjects were tested on the alternate associate that they did not practice, with either the originally trained cue or an independent cue that was also encoded previously. The results were clear: Performing retrieval practice on the dominant sense (e.g., Arm Shoulder) did not impair the later recall of the subordinate sense (e.g., Arm Missile) at all, even though retrieval practice yielded substantial retrieval-based strengthening for the practiced item (see Fig. 3). Practice on the subordinate sense, however, caused retrieval induced forgetting of the dominant sense. Similar results were obtained, regardless of whether subjects were tested on the unpracticed competitor with the homograph (Arm M__) or the independent test cue (e.g., Target-M__ for missile). Thus, even when subjects performed as many as 20 retrieval practice trials on the dominant sense, little retrieval-induced forgetting was observed. Taken together these results argue against associative blocking accounts of retrieval-induced forgetting, but are consistent with idea that this phenomenon depends on the need to override prepotent memories, as would be expected if inhibitory processes are recruited to suppress those memories (see Conway & Engle, 1994 for a related discussion of the role of inhibitory processes in resolving interference in memory span tasks; see also, Lustig, Hasher, & Toney, 2001, for a recent review of work on inhibitory processes in cognitive aging).

Strength independence

Our early work on retrieval-induced forgetting was initially premised on the classical view that strengthening some items would impair later retrieval of other associates (Anderson et al., 1994). However, we quickly discovered that the degree to which practiced items are strengthened does not predict how much retrieval induced forgetting was observed. In fact, as highlighted in the preceding sections, practiced items can be significantly strengthened without causing impairment: Retrieval practice on target items does not impair low taxonomic frequency competitors, subordinate meanings of ambiguous words, nor even high frequency exemplars, provided that retrieval practice is




Fig. 3. Results of a study by Shivde and Anderson (2001): (A) plots the percentage of practiced items (e.g., Arm Shoulder) that subjects recalled on the final test as a function of the number of retrieval practices it received; (B) shows that even after extensive practice on the practiced items, recall of weaker competitors (e.g., Arm Missile) is not impaired on the final recall test (when tested with Arm M__); (C) shows the recall data when the unpracticed competitor was tested with an independent probe (e.g., Target M__ for Missile), again revealing no impairment as a result of practice.

noncompetitive. Nor does inhibition occur if practiced items are strengthened by repeated study exposures instead of retrieval practice. In all of these cases, strengthening occurred, with no resultant inhibition, even when the amount of strengthening was identical to or even greater than that observed in other conditions in which retrieval-induced forgetting was found. When analyses are restricted to cases in which retrieval-induced forgetting is found, the magnitude of the impairment bears little quantitative relationship to the degree of facilitation on practiced items. Together, these findings argue that impairment is independent of the strength of the practiced item.

The property of strength independence is surprising, given the historical emphasis on the role of competition in producing interference (e.g., McGeoch, 1942; Melton & Irwin, 1940; see also, Anderson, 1983; Mensink & Raaijmakers, 1988). However, given the frequent co-occurrence of strengthening and impairment across a variety of paradigms (e.g., retroactive and proactive interference, part-set cuing, list-strength effects), the emphasis on strength as a cause of forgetting makes sense. Anderson et al. (1994) noted however that nearly all paradigms that appear to provide evidence for strength-dependent competition have confounded strengthening with some form of retrieval-induced forgetting. In studies of retroactive interference, for example, one typically cannot disentangle the effects of strengthening word pairs from the second list (e.g., Dog-Sky) from the suppression of first list responses (e.g., Dog-Rock). This ambiguity arises because word pairs from the second list are typically strengthened by repeated study/test cycles, a procedure which conflates strengthening of those pairs with retrieval practice. In list-strength effect studies, one cannot disentangle the effects of strengthening one half of the list of words from the heightened output interference that those strengthened items cause for the remaining nonstrengthened words on later free recall tests. If left to recall items in any order, subjects typically begin with the strengthened items, which is likely to inhibit the remainder of the list. Similar problems occur in part-set cuing studies, which often do not adequately

control for output interference biases (overt or covert) created by providing part-set cues (see Anderson & Neely, 1996; Nickerson, 1984; Roediger & Neely, 1982, for reviews of part-set cuing and related research). Thus, although strengthening some items (even without retrieval practice) appears to impair nonstrengthened competitors, such effects may be better explained in terms of retrieval-induced forgetting.

The difference between the foregoing studies and recent demonstrations of strength independence lies primarily in the attempt to separate the process of strengthening from retrieval-induced forgetting. For example, by using extra study exposures, we have been able to strengthen practiced items without retrieval practice, so that we can see whether the added strength for those items would impair the delayed recall of competitors. Also important has been our effort to control the order in which subjects recall studied items, so as to prevent biases in output order typically created by strengthening manipulations. Towards this end, we have used letter stem cued recall tasks (e.g., Fruit B__) to force subjects to recall nonstrengthened items before strengthened items, and reduce test-time retrieval induced forgetting. By controlling these factors, we have found that strengthening does not cause forgetting of competitors. In a similar vein, Bäuml and colleagues (Bäuml, 1996, 1997, 1998) have also attempted to separate the influences of strengthening and inhibition in other experimental procedures such as the retroactive interference and list strength effect paradigms. Bäuml (1996) found that strengthening a second list of words by increasing study time did not increase retroactive interference on the first study list. Bäuml (1997) showed that the list strength effect virtually disappears if biases in output order are eliminated: strengthening half of a study list through extra study does not impair the later recall of the other list half as long as those nonstrengthened items are tested first in the recall sequence. Thus, a variety of interference effects that have been attributed to strength-dependent competition may arise from the recruitment of inhibitory control processes during retrieval.

Delay dependence?

There is some evidence that retrieval-induced forgetting may recover over time. In one study, MacLeod and Macrae (2001) had subjects perform retrieval practice immediately after encoding had ended. After retrieval practice, subjects were tested either immediately or on the following day. Interestingly, retrieval induced forgetting was observed on the immediate test, but not for those subjects tested after 24 h. In a follow-up study, MacLeod and Macrae (2001) replicated this finding, but also showed that when retrieval practice was performed after a 24 h delay, significant retrieval-induced forgetting was observed on a test given immediately afterwards. Thus, although the impairment may dissipate within 24 h, introducing a long delay between study and retrieval practice did not insulate subjects from retrieval-induced forgetting. In at least some circumstances with some types of materials, the inhibitory effects of retrieval practice recover over time, a finding similar to spontaneous recovery observed in retroactive interference (e.g., Postman, Stark, & Fraser, 1968), directed forgetting (Wheeler, 1995), and the verbal overshadowing paradigm (Finger & Pezdek, 1999). Interestingly, this recovery from inhibition occurs even though practiced items still exhibit significant facilitation after the same delay, again suggesting that differential strength does not cause impairment.

It is not clear, however, whether MacLeod and Macrae's particular delay is needed for people to recover from retrieval-induced forgetting, or even whether recovery *always* occurs. Although MacLeod and Macrae's findings suggest a particular recovery interval, this finding may not generalize to other materials or training protocols. For instance, retrieval induced forgetting may be quite long lasting given different parameters for retrieval practice. Consider learning the new telephone number of a friend whom you call frequently. Initially, their old number will intrude into consciousness when you want to dial their new number. But after dialing the new number enough times over a protracted period, the old number eventually stops intruding. Given enough practice with the new number (over months or a year), one may become completely unable to recall the old telephone number. This inability will likely persist indefinitely, even when you have periods during which you do not call your friend (can you remember your old phone number 3 residences ago?). This suggests that if retrieval practice occurs frequently and is distributed over long time periods, inhibition effects may be long-lasting, although this at present remains an empirical issue.

Theoretically, recovery need not occur at all, however, even if inhibitory processes produce retrieval-induced forgetting. In fact, impairment of nearly any duration may be possible, depending on the mechanisms by which inhibitory effects produce memory failure (Anderson & Spellman, 1995). One can separate the theoretical mechanism that induces retrieval-induced

forgetting from the one that sustains it. For example, inhibitory processes may initially deactivate the units involved in representing a competing memory trace, and this deactivation may be short-lived. However, this deactivation may itself cause a structural change that persists well beyond the initial period of inhibition. For instance, the constituent features of the trace may become less tightly bound or a consolidation process that might have otherwise been ongoing may be terminated (e.g., Müller & Pilzecker, 1900). Inhibition of a subset of features in a trace may lead to reductions in the associations of those features to others that remain active, via the mechanisms of hebbian learning. In an entirely different approach, inhibition might be sustained tonically, even after retrieval practice has ended. The persisting strength of practiced items might sustain inhibition on competitors, via lateral inhibition.

The foregoing theoretical possibilities are not intended to be strong claims about the mechanisms underlying retrieval induced forgetting. They are described merely to illustrate a theoretically crucial point: not enough is presently known about how inhibition is manifest mechanistically to strongly constrain predictions about whether inhibition should recover over time. Inhibitory theories exist that can be consistent with short and long-lasting inhibition (see Anderson & Spellman, 1995 for similar arguments). Which of these mechanistic approaches to retrieval-induced forgetting provides the best account remains to be established.

Generality

Many of the studies discussed so far have used verbal categories to study retrieval-induced forgetting. However, this phenomenon has now been observed with a variety of stimulus classes. For example, Ciranni and Shimamura (1999) found that when subjects learned the locations of colored objects (e.g., squares, circles, triangles or odd, difficult to name shapes), recalling information about one of the objects (e.g., its color or shape) led subjects to forget properties of other objects with the same shape. Using variations of this procedure, they induced subjects to forget the color, location, and shape of the other objects, and found that this impairment only occurred with retrieval-practice and not with extra study exposures. Studies of fact learning have found that retrieving some facts about a topic impairs recall for other facts about that topic (e.g., Anderson & Bell, 2001; Macrae & MacLeod, 1999; Radvansky, 1999). In fact, retrieving some facts about a topic (e.g., The actor is looking at the tulip) not only impairs other facts that directly compete with it (e.g., The actor is looking at the violin), but also facts that share concepts with the competing facts (e.g., The teacher is lifting the violin), replicating and generalizing the cue-independent impairment observed by Anderson and Spellman (1995). Similar cue-independent impairment occurs in the fan interference

paradigm (Radvansky, 1999; however, see Anderson & Reder, 1999). Koutstaal, Schacter, Johnson, and Galuccio (1999) found that reviewing photographs of novel actions that subjects had performed two days earlier (e.g., “trace the outline of this boomerang”), impaired their later recall of the other actions they had performed. Thus, memory can be impaired for one’s own physical actions. Using an eyewitness memory paradigm, studies have shown that interrogating subjects about some details of a mock crime scene impairs memory for other related details (MacLeod, 2002; Shaw, Bjork, & Handal, 1995), a finding that may have significant practical applications. In a related vein, recent work has shown that the tendency for people’s memory for an eyewitness event to be distorted by misleading post-event information may rely to some degree on retrieval-induced forgetting (Saunders & MacLeod, 2002). Saunders and MacLeod found that people were far more likely to inappropriately remember misinformation on a later test when they had earlier performed retrieval practice on other aspects of the event for which the misinformation was introduced. This suggests that vulnerability to misinformation acceptance is heightened when access to the original memory is weakened by inhibition.

Implications of retrieval induced forgetting for social psychological phenomena have also been explored. For instance, Macrae and MacLeod (1999) demonstrated that recalling some traits of a person impairs the retrieval of their other personality traits later. Dunn and Spellman (2003) recently demonstrated that when people repeatedly retrieve individuating traits of a person about whom they recently learned, stereotypic traits of that person were inhibited. Interestingly, the tendency for stereotypic traits to be inhibited was moderated by subjects’ prior belief in the stereotype: subjects who were more prone to believe in the stereotype showed greater resistance to inhibition. Macrae and Roseveare (2002) found that self-relevant encoding may also render information resistant to inhibitory effects. Subjects were presented with a list of words and told that the items on the list were gifts that were purchased. Some subjects were asked to imagine that these were gifts that they had purchased themselves; other subjects were asked to imagine that the gifts were purchased by a best friend or by an unspecified other. Following this encoding phase, the standard phases of the retrieval-practice paradigm were employed. Macrae and Roseveare found that when subjects imaged purchasing the gifts themselves (self relevant encoding), retrieval-induced forgetting was completely eliminated, whereas the inhibition effect remained robust in the other encoding conditions. The protective effect of self-relevant encoding may be an instance of the protective effects of integration (see later section “Integration as a moderating factor”).

Some evidence suggests that retrieval impairs recognition memory for perceptual experiences. Research on

verbal overshadowing has shown that describing a recently viewed face impairs later recognition of that face (Schooler & Engstler-Schooler, 1990). Similarly, describing the flavor of a wine impairs its later recognition (Melcher & Schooler, 1996). Schooler, Fiore, and Brandimonte (1997) suggested that describing a perceptual memory may be a form of retrieval practice. Subjects may retrieve those aspects of the percept easiest to verbalize, eschewing other information that, while not as easy to describe, is crucial to recognition. For example, subjects might focus on verbalizable characteristics of a face, such as basic features (nose, mouth), rather than configural information about the spacing of features in relation to each other. Selectively retrieving features may suppress configural information. Although verbal overshadowing has also been found when the need for retrieval is eliminated (e.g., when subjects receive a description generated by another person), those effects go away if subjects are asked to base their recognition judgments solely on memory for the photograph and not the verbal description. In contrast, subjects generating their own descriptions are not helped by such instructions (Dodson, Johnson, & Schooler, 1997; Meissner & Brigham, 2001). These findings suggest that active retrieval is important to generating a robust effect, just as with retrieval-induced forgetting. Dodson et al. also found that describing another face impaired recognition for the one originally studied, indicating a generalized suppression of face memories (see Anderson & Spellman, 1995, for a potentially related finding termed “second-order inhibition”). Thus, retrieving specific features of a perceptual memory while describing it may inhibit other aspects of the memory.

If inhibitory control mechanisms resolve interference in memory retrieval generally, we would also expect them to be at work in semantic memory. Consistent with this, Bäuml’s (2002) found that episodic memory for several studied exemplars of a category was impaired if subjects generated new exemplars of the same category from semantic memory during the interval between study and test. However, episodic recall was unimpaired when this “semantic generation practice” was replaced by study exposures of the same novel exemplars, showing that impairment derived specifically from semantic retrieval. In a related study, Blaxton and Neely (1983) found that subjects were slower to generate a critical target exemplar (Fruit A___) from semantic memory after they had generated four other prime exemplars from that same category. In contrast, subjects were faster to generate the same target when the prime items were presented intact to subjects for speeded naming. In recent work, Johnson and Anderson (in press) have shown that repeatedly generating associates to the subordinate verb meaning of a homograph from general knowledge (e.g., Prune T_ _M for Prune Trim) reduced the availability of its dominant noun meaning, as

measured by an independent probe test in which subjects free associated to novel test cues (e.g., Yogurt F___ for “fruit,” which is related to the noun sense of Prune). This finding builds on work on lexical ambiguity resolution suggesting that the contextually inappropriate sense of a homograph may be suppressed (Gernsbacher & Faust, 1991; Simpson & Kang, 1994), by establishing that impaired access to homograph meanings is cue-independent. Parallel findings have been observed in episodic memory experiments using homographs, which establish that inhibitory effects are recall specific, consistent with properties of retrieval-induced forgetting more generally (Shivde & Anderson, 2001). Finally, research using the rare-word paradigm has found that difficult semantic retrievals recruit inhibitory processes: When subjects struggle to recall the meaning of an unusual, infrequently encountered word that is weakly represented in memory, related concepts appear to be impaired (Barnhardt, Glisky, Polster, & Elam, 1996; Dagenbach, Carr, & Barnhardt, 1990; see also Thompson-Schill, 1997 for an interesting discussion of the role of the left prefrontal cortex in controlling selective retrieval from semantic memory). Taken together, these results argue that retrieval induced forgetting is not limited to episodic retrieval, or to taxonomic categories; rather, it is a general consequence arising whenever inhibitory mechanisms are recruited to guide selection in the face of competition from distracting memories.

Moderating and masking factors in retrieval-induced forgetting

The preceding review describes evidence showing that whenever we try to selectively retrieve a target item from long-term memory, other competing memories associated to the cue guiding retrieval will be suppressed. Although this is generally true, it is perhaps not surprising that there are factors that can either moderate or mask the effects of inhibition. Moderating factors are those that genuinely alter the magnitude of inhibition that the competitors of main interest suffer during retrieval practice; masking factors are those that alter the later behavioral measure of inhibition without affecting the magnitude of inhibition that actually transpired during retrieval practice. Appreciating these factors is a fundamental part of understanding the behavioral conditions under which inhibitory control leads to forgetting. These factors can be divided broadly into those concerning how memories are represented, how retrieval practice is performed, and how memory is ultimately assessed after inhibition has been induced. We discuss these in turn.

Representational factors that moderate or mask inhibition

When predicting how much inhibition will occur in a given population or in a particular condition, one

cannot consider the characteristics of the inhibitory mechanism in isolation. Inhibitory mechanisms act on memory representations. These representations may vary in type, structure, content, or strength, and these variations may moderate the impact of inhibitory processes or even the necessity of inhibition. Concern over this possibility has a long history in research on interference, and is reflected in classic work on verbal mediation (see Horton & Kjeldergaard, 1961; Jenkins, 1963; Kjeldergaard, 1968; Postman, 1971, for reviews) and similarity effects (see Osgood, 1949, for a review) on retroactive interference, integration effects on fan interference (see, e.g., Radvansky, 1999; Radvansky & Zacks, 1991; Smith, Adams, & Schorr, 1978), and the effects of level of learning on the magnitude of retroactive interference or fan effects that are observed (see, e.g., Postman, 1971 for a review for retroactive interference; see Hayes-Roth, 1977, for a review concerning fan interference). It is thus not surprising that these factors are also involved in moderating retrieval induced forgetting as well. I review the evidence for two such moderating factors: integration and similarity. I also describe a representational factor that masks inhibition—baseline deflation.

Integration as a moderating factor. The amount of retrieval-induced forgetting depends strongly on how well integrated the to-be-retrieved memories are with the practiced competitors. Although there is some variation in how the term integration has been used in the literature, we have used it to refer to the existence of interconnections between items sharing a common retrieval cue—connections formed either on the basis of pre-experimental relationships, or novel interrelationships discovered during the course of an experiment. For instance, suppose subjects studied “Animals” such as Deer, Dog, Bear, Canary, Goat, and Cow. In addition to studying these items in relation to their shared category label, subjects might form inter-relationships between items such as Deer and Bear (Wild things that you hunt), Goat and Cow (farm animals), Dog and Canary (pets), or Dog and Deer (an image of a dog chasing a deer). These inter-relationships could be based on semantic similarity (e.g., Dog, Wolf), associative relatedness (Dog Bone), or even on more elaborate encoding of relations (e.g., interactive imagery).

In general, when subjects integrate the associates of a cue, it insulates nonpracticed exemplars from retrieval-induced forgetting (Anderson & McCulloch, 1999). Anderson and McCulloch demonstrated this using the retrieval-induced forgetting design of Anderson et al. (1994), but with one change: at the time of encoding, subjects either were or were not encouraged to find inter-relationships among the exemplars of a category. Subjects who were asked to integrate exemplars showed a significant reduction in retrieval induced forgetting (and in some cases, it was completely eliminated).

Interestingly, even some subjects who were not asked to integrate reported that they had done so on their own, as measured by a post-experimental questionnaire. These subjects showed the same reduction in retrieval-induced forgetting as the group instructed to use integration. The more study time subjects were allowed, the higher the reported use of spontaneous integration strategies, and the lower the amount of retrieval induced forgetting. Similar integration effects were observed in retrieval-induced forgetting experiments using propositional materials (Anderson & Bell, 2001): when subjects reported integrating multiple facts about a topic (e.g., The ant crawled on the rock, The ant crawled on the table), retrieval-induced forgetting was either reduced or eliminated. These latter effects were even found in an incidental encoding task in which subjects were asked to form vivid mental images of the situations represented by the sentence; when subjects reported incorporating multiple facts into a single image, inhibition was significantly reduced. Thus, although retrieving some associates of a retrieval cue generally impairs other associates that become activated in the process, integration poses a strong boundary condition on when this impairment occurs.

Similarity as a moderating factor. The amount of inhibition that retrieval will cause also depends on semantic similarity between the associates of a cue. The nature of this relationship is complex, however, as illustrated by the studies of Smith and Hunt (2000) and Bäuml and Hartinger (2002). Smith and Hunt (2000) adapted the retrieval practice procedure so that the degree of within-category similarity might be varied. Specifically, they altered the study phase to encourage the encoding of either similarities or differences between exemplars of a category. For the similarity encoding group, subjects viewed all six exemplars of the category at once and were asked to find a way that the item at the top of the list was similar to all of the remaining items. Shared features were then generated in turn for the other five exemplars. After encoding the categories in this way, subjects went through the remaining phases of the retrieval practice procedure. The difference encoding group followed the same steps, but was asked instead to find one feature that made the top item different from all of the remaining items. Smith and Hunt found that encoding differences between exemplars abolished retrieval induced forgetting, but encoding similarities yielded robust impairment. They concluded that inter-item similarity increases retrieval-induced forgetting.

However, Bäuml and Hartinger (2002) found a pattern that appears to directly contradict that observed by Smith and Hunt. These authors also sought to manipulate the inter-item similarity between the exemplars of a category, but they manipulated similarity by varying whether or not the unpracticed competitors in a category (e.g., Fruit Lemon) were drawn from the same

subcategory (e.g., Citrus) as the to-be-practiced items (e.g., Fruit Orange), or a different subcategory (e.g., Fruit Cherry). After the encoding phase, subjects engaged in retrieval practice in the usual fashion, and then were given a final category-plus-letter stem cued recall test. In contrast to Smith and Hunt (2000), Bäuml and Hartinger found that increasing interitem similarity reduced retrieval-induced forgetting. Bäuml and Hartinger replicated this pattern using an output interference paradigm instead of retrieval practice, and argued that similar mechanisms underlie these two phenomena.

The Smith and Hunt (2000) and Bäuml and Hartinger (2002) findings are not necessarily contradictory. In recent work, Anderson et al. (2000) explored whether semantic similarity might have different effects on retrieval induced forgetting depending on whether one is concerned with what they termed *target-competitor similarity* or *competitor-competitor similarity*. As illustrated in Figs. 4A and B, the unpracticed competitors in a category undergoing retrieval practice can either vary in (a) how similar they are to the target items receiving retrieval practice (i.e., target-competitor similarity), or (b) how similar they are to each other, independent of how similar they may be to the retrieval practice targets (i.e., competitor-competitor similarity). Anderson et al. (2000) argued that these two dimensions should have very different effects on retrieval-induced forgetting, based on the distributed representation approach proposed by Anderson and Spellman (1995). According to this approach, increasing target-competitor similarity from a moderate level (top of Fig. 4A) to a very high level (bottom of Fig. 4A) should diminish retrieval induced forgetting. Less impairment should be observed because, according to the model, the recall probability of an item reflects the summed activation of all of its features. Because high target-competitor similarity leads many of a competitor's features to overlap with the retrieval practice target, the facilitatory effects of retrieval practice on shared features will compensate for or possibly even outweigh the inhibition suffered by the competitor's distinctive features. On the other hand, increasing competitor-competitor similarity from a moderate level (top of Fig. 4B) to a high level (bottom of Fig. 4B) should magnify the amount of retrieval-induced forgetting. More impairment should be observed because in the high similarity condition, the impact of suppressing a single feature that overlaps two different exemplars will be realized through the impairment of two items, not just one; thus, the behavioral effect of applying the inhibition to highly overlapping representations will be exaggerated, even if the same amount of inhibition is applied.

Anderson et al. (2000) tested these hypotheses by separately manipulating the degree of target-competitor and competitor-competitor similarity. Following Smith and Hunt, they held the study materials constant and

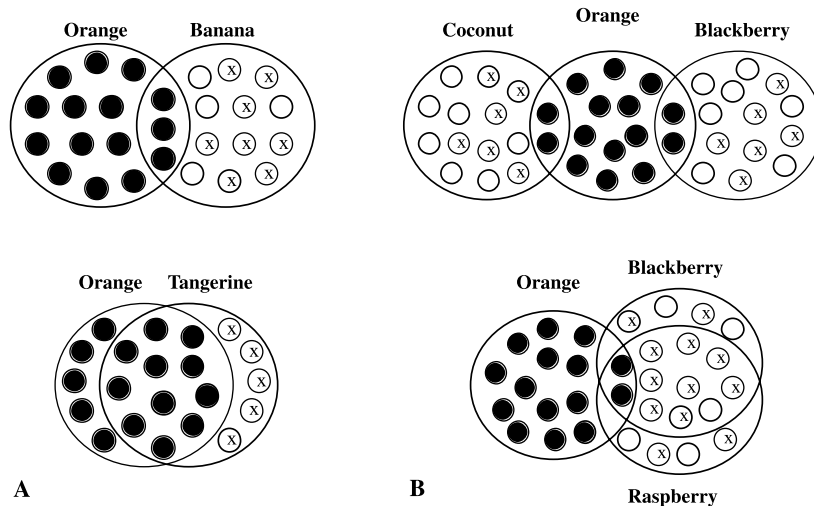


Fig. 4. Illustration of two different dimensions of similarity, as distinguished in the two-factor model of Anderson, Green, and McCulloch (2000). Memory items (larger circles) are represented here as sets of semantic features (small circles). Similar items overlap in feature space (as represented by overlapping larger circles). Retrieval practice is assumed to increase the activation of practiced features (darkened circles) and to inhibit some of the features of the competing, similar pattern (lighter circles with Xes in them), but not other features of the competitor (small white circles). (A) An illustration of how target–competitor similarity (similarity between the practiced item and an unpracticed competitor) can be low (top half) or high (bottom half) and how this influences inhibition. With high target–competitor similarity, a greater proportion of a competitor’s features overlap with the practiced item and are strengthened, compensating for inhibition on the remaining features. (B) An illustration of how competitor–competitor similarity (similarity amongst the competitors themselves) can be low (top) or high (bottom), and how this may influence inhibition. With high competitor–competitor similarity, inhibiting the same number of units has a greater impact on the two competitors, because the inhibition affects features shared by the two items.

manipulated similarity by asking subjects to identify similarities or differences between exemplars. However, instead of having subjects do this for all pairwise comparisons within a category, subjects were presented with either target–competitor or competitor–competitor pairings, to control the dimension of similarity that was manipulated. Following this similarity (or difference) encoding phase, the remaining steps of the retrieval practice procedure were done in the typical fashion. The results were striking: in the target–competitor condition, significantly less inhibition was found when subjects were asked to find similarities than when they were asked to find differences between items during encoding. In fact, subjects who were asked to find target–competitor similarities showed significant retrieval-induced facilitation of competing items, not inhibition. In the competitor–competitor condition, however, more inhibition was found when subjects were asked to encode similarities than when they were asked to encode differences. Indeed, the difference encoding condition yielded no significant inhibition. These findings strongly support the idea that competitor–competitor similarity has an opposite effect on inhibition than target–competitor similarity, as suggested by the Anderson and Spellman (1995) distributed approach. Anderson et al. (2000) argued that these findings help to reconcile the

conflicting findings of Bauml and Hartinger and Smith and Hunt, as well as analogous inconsistencies in the literature on the role of similarity in classical interference studies (see Anderson et al., 2000 for a discussion).

Baseline deflation as a masking factor

When considering how representational variables that might affect inhibition, it is also important to attend to the representation of baseline items. Retrieval-induced forgetting may be masked if the baseline used to assess inhibition is also affected by retrieval practice. Such “baseline deflation” may arise in two ways. First, as Anderson et al. (1994) noted, practiced and baseline categories are represented in a common episodic context. Retrieval practice may therefore suppress items in baseline categories because they share contextual features with items undergoing retrieval practice. To the extent that baseline categories are also suppressed by inhibitory processes, the ability to determine how much inhibition has taken place on within-category competing exemplars is compromised. This possibility is arguably consistent with several findings in the output interference and retrieval practice literatures. For example, recall probability declines for categories or paired associates that are tested later in a testing sequence, even when those categories or paired associates are not

apparently similar and do not explicitly share any cues (e.g., Roediger & Schmidt, 1980; Smith, 1971). Furthermore, Tsukimoto and Kawaguchi (2001) found that baseline categories can be impaired by retrieval practice, as compared to performance on those same baseline categories in a control group who did not perform retrieval practice. These findings suggest that shared contextual features may precipitate baseline suppression, reducing measured inhibition. This factor thus masks inhibition because it affects estimates of the amount of inhibition on competing items without altering the true level of inhibition that actually took place for those items.

A second source of baseline deflation can arise when baseline and practiced categories are similar. To the extent that baseline categories share semantic features with items that are inhibited by retrieval practice, retrieval-induced forgetting may generalize to those categories. Anderson and Spellman's (1995) cross category inhibition findings provide a case in point: Practicing red items such as Red-Blood not only suppressed items explicitly studied under the Red category (e.g., Red Tomato), but also other red items that were studied and tested under a separate category (e.g., Food Radish). Even non-red Food items (e.g., Food Bread) were inhibited after subjects practiced Red-Blood, suggesting that the inhibition of items that directly competed with the retrieval practice target (e.g., Red-Tomato) semantically generalized to other items that overlapped with them in semantic features. In a similar vein, Anderson and Bell (2001) found that practicing some facts about a topic (e.g., The actor is playing the guitar) impaired not only other facts sharing that topic (e.g., The actor is playing the oboe), but also facts studied under a different topic but sharing the same relation and category (e.g., The teacher is playing the drum). Thus, impairment generalized across topics, based on semantic similarity. Anderson and Bell (2001) were able to measure the inhibition of the latter items because they included additional baseline topics that did not share the same relation and category with practiced items (e.g., The box is in the warehouse, The mop is in the pub). The generalized suppression was circumscribed to items with specific overlap in semantic relations with items studied with the practiced topic, and could not have been produced by a global contextual similarity. These findings strongly suggest that studies of retrieval induced forgetting need to take appropriate measures to ensure that within-subjects baseline conditions are as dissimilar from practiced categories as possible. Baseline suppression effects such as these might be one reason why Anderson and Reder (1999) failed to find evidence for cue-independent impairment in their fan effect paradigm: All of their propositions were constructed using the same semantic relation and object class (all were "is in" facts, such as "The lawyer was in the park").

Retrieval-practice factors that moderate inhibition

The amount of inhibition that occurs is also likely to depend on the amount of attention given to the different cues provided for retrieval practice. In particular, any type of retrieval practice that minimizes the need to resolve interference between competing items is unlikely to produce inhibition. Consider, for example, the study by Anderson et al. (2000) reviewed earlier. When subjects were given the category and asked to recall the exemplar based on stem cues (e.g., Fruit Or___), significant retrieval-induced forgetting was observed; however, when subjects were given the exemplar, and asked to recall the category (e.g., Fr___ Orange), there was no impairment. This pattern is likely to have arisen because the cue in the latter case—Orange—was associated to the category, but not to other exemplars in the category, eliminating competition that would lead to impairment. Similarly, if subjects were asked to perform retrieval practice without the category label (e.g., Or_n_e for Orange), other exemplars in the category are unlikely to interfere and thus may not be impaired. Subtler cases may also be possible. For example, even when subjects are given the category and a fragment cue for retrieval practice, subjects might focus their attention on the fragment cue—that is, they may solve the retrieval practice task by circumventing interference caused by the shared category. This seems especially likely when the fragment cue is highly informative or draws attention. For example, if multiple letters are provided (e.g., Fruit B_n_n_), subjects might spend more of their time focusing on the distinguishing letter features, trying to solve the fragment by "sounding the word out." In general, any factor that reduces attention given to the shared cue and focuses it on the distinguishing cue is likely to reduce activation of competitors and therefore reduce inhibition.

Test factors that moderate, mask, or exaggerate inhibition

In our initial studies of retrieval-induced forgetting, we measured subjects' final memory performance with a category cued recall test. Subjects were provided with each studied category name in turn, and asked to recall all of the studied exemplars in any order. Inhibition has been found consistently with this type of test (Anderson & Bell, 2001; Anderson et al., 1994; Anderson & McCulloch, 1999; Anderson & Spellman, 1995; Butler, Williams, Zacks, & Maki, 2001; Macrae & MacLeod, 1999; Nader, Coles, Brigidi, & Foa, 2001; Smith & Hunt, 1999), even when the shared cue is not categorical in nature (e.g., Anderson & Bell, 2001; Macrae & MacLeod, 1999). Other tests have also been used, however, not only to characterize the range of conditions under which retrieval-induced forgetting occurs, but also to infer various properties of the effect. In this section, I discuss some of the work that has been done with alternative testing formats, with an emphasis on factors that may moderate or mask inhibitory effects.

Output interference effects. Depending on the test type that one uses, the amount of retrieval-induced forgetting may reflect at least two sources: impairment arising from the earlier retrieval practice phase, and impairment produced by the final recall test. The final test contributes a second source of impairment because the strengthening of practiced items during the earlier retrieval practice phase leads those items to be recalled earlier in the final test sequence. Because this will delay unpracticed competitors until later in the output sequence, these items are subject to additional retrieval induced forgetting—that is, they are subject to exaggerated output interference, relative to baseline categories. This bias in the retrieval of practiced items is interesting because it may provide one mechanism by which the inhibitory effects of retrieval can be reinstated on a recurring basis, even when initial inhibitory effects have dissipated (Anderson & Bell, 2001). However, the contribution of test-based sources of impairment can impair clear theoretical inferences about the conditions producing inhibition, and so it is necessary to consider this factor in assessing inhibition. Such inferential difficulties are most likely to arise in test formats that allow subjects to report items in any order they wish, although they are not restricted to those types of test. Theoretically, exaggerated output interference is neither a moderating nor a masking factor, because it does not alter the amount of inhibition that actually took place during retrieval practice, nor does it prevent us from seeing this effect; it does, however, alter the measured estimate of retrieval-practice based inhibition.

The contribution of output interference is of greatest concern in two varieties of experiment: when one wants to establish the retrieval practice phase as the primary source of impairment, and when one is concerned with variations in the amount of inhibition that have occurred across different conditions or groups. Knowing whether inhibition primarily reflects events in the practice phase is important, for example, in determining whether extra study exposures cause inhibition. Retrieval practice and extra study exposures both strengthen the practiced items, so that on a delayed recall test, those items are likely to be recalled early in the recall sequence. If subjects are free to recall items in any order, unpracticed competitors in both of these conditions will be subject to greater output interference (test based retrieval-induced forgetting) than corresponding items in baseline categories. Thus, even if extra study exposures produced no inhibition during the practice phase, significant impairment might be observed on the final test, leading one to conclude that extra study exposures caused inhibition. Similarly, if one is concerned with how long retrieval-induced forgetting lasts, one must be sure that test-based sources of impairment do not contribute to the measure of inhibition, or one might be led to believe that inhibition lasts longer than it truly

does. Finally, when one wants to compare the relative amount of inhibition across two conditions or groups, it is essential to disentangle different sources of inhibition. Greater inhibition may occur in one condition, for example, merely because greater associative strengthening for practiced items produced greater output-based effects in that condition.

The most straightforward way to distinguish the practice and test-based contributions to retrieval-induced forgetting is to use a test such as category-plus-stem cued recall that enables one to control recall order (Anderson et al., 1994). In a typical study, the final test is composed of trials in which each exemplar is cued with its category name and a one letter stem for the exemplar. Importantly, subjects are first cued to recall all of the unpracticed items from a category, then all of the practiced items, or vice versa. Comparisons are then made to baseline items tested in the corresponding halves of their respective categories. It is typically assumed that recall impairment observed when all unpracticed exemplars are tested before practiced items must reflect the lingering effects of the retrieval practice phase, for the simple reason that practiced items have yet to be recalled. When unpracticed exemplars are tested first in this way, significant retrieval-induced forgetting is typically found (Anderson & Bell, 2001; Anderson et al., 1994; Anderson et al., 2000; Anderson et al., 2000; Anderson & McCulloch, 1999; Bäuml, 2002; Bäuml & Hartinger, 2002), although sometimes it is reduced in magnitude from the effects observed with category cued recall without letter stems. This finding makes sense given the elimination of output interference from the effect. In comparing the recall of items tested in the first half of their categories to those tested in the second half, output interference is typically observed on this kind of test, reinforcing the importance of isolating the two sources of impairment. By using this type of testing procedure, several studies have found that strengthening competitors does not reliably impair related items when output interference is controlled (for retrieval-induced forgetting, see Anderson et al., 2000; for list-strength effects, see Bäuml, 1997; see also Bäuml, 1996 for a conceptually similar finding for retroactive interference).

However, using category-plus-stem cued recall is not sufficient to ensure that output interference has been adequately matched across baseline and practiced categories. There are cases in which output interference differences can arise even when recall order is fixed. In particular, category-plus stem cued recall tests in which the practiced and unpracticed exemplars of a category are randomly interspersed in the recall order do not adequately control for output interference. For example, suppose that subjects study the items Orange, Banana, Lemon, Cherry, Apple, and Grape as members of the Fruits category, and then perform retrieval practice on

Orange Banana and Lemon. On the final recall test, the recall sequence Grape, Cherry and Apple, Banana Orange Lemon would control for output interference biases, whereas Apple, Orange, Cherry, Banana, Grape, and Lemon would not. Although the latter format tests items in a fixed order that is constant across practiced and baseline categories, output interference is not matched. This is because when the category is practiced, there is a much greater likelihood of recalling the practiced items (Orange, Banana, and Lemon) than the items in the corresponding positions for the baseline category. Thus, more test-based output interference will be exerted on unpracticed competitors in the practiced category when practiced and unpracticed items are interspersed (at least on those that follow practiced items). If it is important to ensure that inhibition effects are not being produced at the time of output, weaker items should be tested before strengthened items.

Cue priming as a masking factor. When subjects perform retrieval practice, they are typically presented with a category name and the first two letters of the exemplar that they are to retrieve. If retrieval is successful, the practiced item is facilitated, and competing items are suppressed. However, retrieval practice introduces another factor as well. Given that the practiced category is typically presented nine times in the standard retrieval practice session (3 exemplars are practiced three times each), the category name enjoys a substantial boost in accessibility. In some circumstances, this cue priming can reduce the amount of inhibition that is measured, without actually influencing the level of inhibition that takes place.

The effect of cue priming on measures of inhibition can be seen in our first experiment on retrieval-induced forgetting (Anderson, 1989). This experiment employed the basic retrieval practice paradigm, except that we used free recall as our final test instead of category cued recall. The results can be seen in the left side of Fig. 5A

for the practiced, competitor, and baseline conditions. As might be expected, retrieval practice facilitated the delayed recall of practiced items; more surprising, however, retrieval practice also facilitated unpracticed competitors, relative to items in unpracticed categories. A comparable result can be seen in the right side of Fig. 5A, which depicts the findings of a highly similar but independent experiment by Smith and Hunt (2000).

The findings of Anderson (1989) and Hunt and Smith (1998) appear inconsistent with the notion that retrieval practice suppresses related items, as has been argued throughout. Indeed, from a behavioral standpoint, these results indicate that under some testing conditions, retrieval practice enhances the recall of related items. However, to conclude that no inhibition occurred in these experiments would be a mistake. Both experiments used a testing format (free recall) that permits cue priming to influence how much inhibition is measured. Specifically, with categorized word lists (and organized lists in general), it is widely believed that subjects adopt a hierarchical retrieval scheme for recalling study items; retrieval progresses first from a representation of the episodic context in which items are studied, to the categories on the list, and next from the category representations to the particular exemplars (see, e.g., Rundus, 1973). Given this multi-stage process, the likelihood of recalling an exemplar is influenced by two probabilities: the probability of recalling the category label, given the context as a cue, and the probability of recalling the exemplar, given that the category label has been recalled. The combination of these factors determines how well practiced items, unpracticed competitors, and baseline items will be recalled. Ordinarily when category cued recall is used, the probability of recalling the category labels is constant at 1.0, because the labels are provided. However, on free recall tests, biases in category recall across conditions become an issue, particularly when more than just a few categories are used and subjects




Fig. 5. Examples of cue priming effects in free recall in studies by Anderson (1989) and Hunt and Smith (1998). Subjects underwent the standard retrieval practice procedure of Anderson et al. (1994) and were tested with free recall instead of category-cued recall. (A) Percentage of practiced items, unpracticed competitors and baseline items recalled on the final free recall test. Practice facilitated the practiced items as well as the unpracticed competitors. (B) The same data as in (A), counting only those items from categories for which subjects recalled at least one item (ensuring category access). Conditionalizing recall in this way reveals a significant retrieval induced forgetting effect that had been masked by primed access to category labels.

may thus forget categories. Given that practiced category labels (e.g., Fruits) are primed, they are quite salient to subjects, leaving baseline categories at a recall disadvantage. The end result is that subjects are more likely to forget whole baseline categories, and miss the opportunity to recall exemplars from those categories. Thus, suppression of unpracticed competitors by retrieval practice may be masked by the overall advantage in the accessibility of practiced categories.

To illustrate how cue priming masked a true deficit in exemplar access in the Anderson (1989) and Hunt and Smith (1998) studies, the free recall data were reanalyzed to focus on only those categories from which subjects recalled at least one exemplar. Our assumption was that subjects who recalled at least one exemplar from a category had accessed the category. Restricting the analysis to those categories would allow us to examine exemplar access for the practiced and baseline conditions, uncontaminated by differences in category access. As can be seen in Fig. 5B, this analysis revealed a pattern of retrieval induced forgetting quite similar to what is ordinarily observed, with unpracticed competitors being recalled more poorly than baseline items. Additional analyses confirmed that the probability of forgetting whole categories (i.e., category “dropout”) was much higher for baseline categories (20%) than it was for practiced categories (1%). Hunt and Smith (1998) found a very similar pattern, as can be seen in the right side of Fig. 5B. These findings illustrate how cue priming can mask inhibition effects when a multi-stage recall process is likely, as it is on free recall tests.

However, cue priming effects are not limited to free recall, nor to categorized word lists. Consider the study of propositional retrieval-induced forgetting by Anderson and Bell (2001). When subjects performed retrieval practice on previously learned facts such as “The ant crawled on the rock,” the later recall of other facts sharing that topic such as “The ant crawled on the table,” was impaired relative to baseline facts such as “The actor looked at the painting.” However, Anderson and Bell cued subjects on their final test with the topic and the relation (e.g., “The ant crawled on the _____,” and “The actor looked at the _____”), sometimes together with a letter stem. If we had instead simply given subjects the cue “The ant” and “The actor,” the final test would likely have become a multi-stage recall test, even though free recall was not used. This is because we used many different topics with different semantic relationships (e.g., is crawling on, is looking at, is in, is eating, etc), most of which could be paired with any topic and so could not be easily guessed. Given only the topic as a cue, subjects would have had to recall the activity or relationship that the topic was engaged in first, followed by the objects of that activity. Because subjects practiced three facts for each practiced topic three times each (e.g., three things that the ant crawled on), the semantic

relationships for the practiced topic would have been far more accessible than the semantic relationship used in the baseline topics. This suggests that even if the object of the unpracticed competitor “The ant crawled on the chair” was suppressed during the practice of “the ant crawled on the rock,” this suppression would be masked by heightened availability of the shared relational concept unless it was provided as a final test cue.

Thus, the complexities introduced by cue priming are not limited to categorized word lists or to free recall. Indeed, cue priming may even influence simple paired associates tests to the extent that subjects link the stimulus and response members with a relation (e.g., when encoding the pair Bird Worm, the relation “eats” is likely to be inferred) that may be forgotten independently of memory for the response. This factor makes it crucial to consider whether the structure of the materials used in a given paradigm, when coupled with the test format, might mask inhibition through cue priming.

Masking through transfer-inappropriate testing effects. Whether inhibitory effects will be observed should depend on the degree to which the memory trace tapped by the retrieval test matches the trace that was inhibited by retrieval practice. To illustrate this, suppose that a subject encodes the pairs Tree-Prune, Tree-Rock, and Trim Prune and then does retrieval practice on Tree-Rock. Later on, suppose that subjects’ memory for Prune is tested either by cuing with Tree-P___ or Trim P___. If retrieval-induced forgetting is found with Tree-P___, but not Trim P___, would it mean that impairment is cue-dependent? If so, would it mean that “Prune” was never inhibited? At first glance, it might seem so, to the extent that cue-independence is an essential feature of inhibition. After all, these tests vary in the cues that they present to subjects, so if impairment depends on which cues are used, it must obviously be cue dependent. This would appear to contradict the property of cue-independence. However, this conclusion does not necessarily follow.

The problem is that the foregoing argument fails to consider the distinction between the *nominal form* of a stimulus, and its *functional representation* by subjects. Although from the standpoint of the experimenter, the word Prune is identical when presented in the pairs Tree-Prune and Trim-Prune, the underlying representations formed by subjects may not be. When studying Tree-Prune, subjects might have encoded prune’s fruit sense, but when studying trim-prune, they certainly would encode its verb meaning instead. When retrieval practice was performed on Tree-Rock, an episodic representation including the fruit sense of prune may have been suppressed, making it less accessible when tested with “Tree P___.” Such inhibition would not be expected to materialize on the test Trim P___, however, because this test taps subjects’ episodic memory for an entirely unrelated concept that was never inhibited (a better

independent probe, in this circumstance, would have been Fruit P___). In essence, the independent probe Trim P___ is simply not testing the same episodic representation that was inhibited, even though it may seem the same from the experimenter's standpoint. This example illustrates how the cue-independence property of inhibition pertains to the particular *functional representation* that is formed by the subject: given that a representation is inhibited, its recall should be impaired, and this impairment should be observable from a variety of cues that tap *that particular representation*. For these reasons, when designing tests to determine whether or not inhibition is present, it is essential to ensure that the test might reasonably be expected to tap the representation that was inhibited by the subject. If not, *transfer-inappropriate testing* may mask the inhibition that actually occurred.

Transfer-inappropriate testing effects may not be limited to stimuli that have different meanings, or to the use of the independent probe method. These effects may also arise when multiple *levels* of representation are possible. For instance, during word encoding, orthographic, phonological, and conceptual representations may each be formed (depending on the orienting task), and these representations may be functionally and anatomically distinct (see Balota, 1994, for a review; see also Roskies, Fiez, Balota, Raichle, & Petersen, 2001 for a discussion of anatomical localization of these different linguistic codes). If different levels of representation are formed for the same nominal verbal stimulus, there is potential for transfer-inappropriate testing to attenuate or mask inhibition. To see this, suppose that performing retrieval practice using a categorically driven cued-recall test such as Fruits Or___ (for Fruits Orange), inhibits conceptually based episodic representations of competing fruits such as Banana. If this conceptually based representation is structurally distinct from the phonological and orthographic representations formed during the initial processing of Banana, little inhibition would be expected for Banana on orthographic or phonologically oriented tests. Retrieval may simply fail to make contact with the representation that was inhibited. The underlying principle behind this idea receives some support from findings in the levels of processing literature: Manipulating the level of processing of words at encoding has dramatic effects on later recall and recognition tests, but, these effects can disappear or even reverse when the final explicit memory test focuses subjects on the lexical and phonemic properties of words (e.g., Fisher & Craik, 1977; McDaniel, Friedman, & Bourne, 1978; Morris, Bransford, & Franks, 1977). On perceptually driven implicit memory tasks such as word fragment completion, word stem completion, and perceptual identification, levels of processing has little effect (e.g., Jacoby & Dallas, 1981; Roediger, Weldon, Stadler, & Rieger, 1992; see Roediger & McDermott, 1993, for a

review). If orthographic or phonological tests are less sensitive to *increases* in the accessibility of conceptually coded information about a word (as is shown by levels of processing dissociations), it seems possible that they might also be less sensitive to *decreases* in the accessibility of those codes produced by inhibition.

The idea that categorically driven retrieval practice primarily inhibits conceptual as opposed to orthographic or phonological levels of analysis receives some support from a recent study by Butler et al. (2001). These investigators employed the retrieval practice paradigm, but varied the nature of the final recall test. Different groups were tested with the standard category cued recall test (e.g., presentation of the category "Bird"), or with one of several lexically oriented implicit and explicit recall tests such as word fragment completion (e.g., cuing subjects with _p_r_ow for the word "Sparrow" with a free completion instruction), word fragment cued recall (e.g., _p_r_ow with an explicit recall instruction), category-plus-fragment cued recall (e.g., Bird, _p_r_ow) or category-plus stem cued recall (e.g., Bird Sp_____). With the exception of category cued recall, these tests focus subjects' attention to varying degrees on the orthographic and phonological features of the cued words. Subjects are likely to complete the fragment _p_r_ow not primarily through conceptually driven episodic recall, but by "sounding out" the answer based on general knowledge of word forms. If so, retrieval-induced forgetting may be attenuated because the test weights a level of representation different from the one that is inhibited. Consistent with this, Butler et al. found no retrieval-induced forgetting on any tests involving letter cuing. These results are compatible with the idea that retrieval-induced forgetting primarily affects conceptually based representations. Unfortunately, it is difficult to attribute these effects to transfer-inappropriate testing because Butler et al.'s experiments are likely to be contaminated by integration strategies during encoding. Subjects were given 8 s to study each exemplar instead of the usual 4–5 s, a procedure likely to increase integration (Anderson & Bell, 2001; Anderson & McCulloch, 1999). This seems especially plausible, given the unusually small amount of retrieval-induced forgetting that they found in their category cued recall condition (5%, compared to the typical 9–20%).

The notion of transfer inappropriate testing is particularly important to consider in connection with experiments examining whether retrieval-induced forgetting affects performance on implicit memory tests. One might argue that if retrieval practice truly inhibits competitors, effects should be observed on indirect memory tests. Caution is warranted here, however, because not all indirect memory tests are the same. Many of the most common tests are perceptually oriented, such as word fragment completion, lexical decision

(some varieties) and perceptual identification. To the extent that these tests tap perceptually based representations, they would not be expected to yield evidence for inhibition, regardless of their implicit/explicit status. A better strategy would be to use conceptually driven indirect tests such as free association, semantic fluency and perhaps category verification, which would be sensitive to variations in the accessibility of a semantic representation. Consistent with this possibility, recent studies have found retrieval-induced forgetting on conceptually driven, but not perceptually driven implicit tests (Perfect, Moulin, Conway, & Perry, 2002; see also Moulin et al., 2002 for further evidence of impairment on a conceptual implicit tests).

The foregoing discussion does not imply that perceptually oriented memory representations cannot be inhibited by retrieval practice. Indeed, the type of representation affected by inhibition should be driven by which representations cause competition during retrieval. This should be determined in part by the nature of the cues guiding retrieval practice, and by subjects' retrieval goals. If the subject is asked to retrieve a studied word that begins with the letters *Ac*, orthographically similar competitors may be more inhibited than semantically related competitors. Although this has not been tested, related research on implicit memory is consistent with this possibility. For instance, Rajaram, Srinivas, and Travers (2001) found that the amount of repetition priming exhibited for a word on either a word fragment or stem completion test was significantly reduced when subjects had to ignore that word's identity during encoding. When subjects were presented with a word colored in red, blue, green, or yellow, and asked to quickly identify the color of the word, subjects exhibited less priming than when they simply had to name the word itself. Although one might attribute reduced priming to reduced encoding in the color naming condition, Rajaram et al. established that the words had been identified sufficiently to cause competition with color naming; the reaction time to name the color of a word was significantly slower than the time to name a neutral stimulus (e.g., a row of *Xes*). Rajaram et al. argued that the diminished repetition priming reflects the inhibition of the word itself, driven by the need to focus attention on the color attribute of the word during the color naming trial—a process they refer to as deselection. If correct, this view suggests that retrieval driven by one perceptual attribute of a stimulus (e.g., color) may suppress other perceptual aspects of that stimulus that cause interference (e.g., visual word form). This effect may later be observed on a perceptually driven implicit test that relies on the rejected attribute. Analogous dynamics may be partially responsible for certain cases of implicit memory blocks driven by orthography of a word (Smith & Tindell, 1997). The standard retrieval-practice experiment with categorically driven

retrieval practice may only reveal inhibition on conceptually driven tests because retrieval practice is conceptually oriented.

Even if conceptually driven implicit tests did not show inhibition, however, it wouldn't by itself imply that inhibition effects did not occur. Here again, it remains possible that the lack of impairment on the implicit test may be due to transfer-inappropriate testing effects. Theoretically, it seems reasonable to distinguish between a general semantic representation of an item (e.g., *Banana*) and an episodic representation of that item as it appeared on a study list. The episodic representation of the item may be composed not only of distinctive contextual features, but also instantiations of semantic features generally used to represent the item in semantic memory. To the extent that such an episodic representation is at least partially structurally distinct from the general semantic representation of the item (the episode-specific component residing perhaps as a bound set of features in the hippocampus, as opposed to neocortex; see, e.g., Norman & O'Reilly, in press), we must consider the possibility that the episode can be suppressed without affecting the general concept of *Banana* (Anderson & Bell, 2001). This form of episode-specific inhibition may be particularly likely when episodic retrieval practice is performed, as in most studies of retrieval-induced forgetting; because retrieval practice is guided not only by a category and a letter stem, but also by a contextual representation of the study list, the episodic representation of a competing item may be the primary source of competition, not the semantic representation of an item. In fact, research has demonstrated that episodic representations can indeed be inhibited: Ciranni and Shimamura (1999) found evidence that novel visuo-spatial representations can be inhibited by retrieval practice, even though these representations clearly do not have well learned semantic counterparts. It may therefore be possible to observe episode specific inhibition in more traditional retrieval-induced forgetting experiments in which the materials also happen to have a corresponding representation in semantic memory. If episode specific inhibition is possible, such inhibition effects should generalize to independent retrieval cues used to test accessibility of that episode (on an explicit test), even when effects do not appear on implicit tests.

Although the Ciranni and Shimamura findings indicate that episode-specific inhibition may occur, a number of considerations suggest that this may not provide a general account of retrieval-induced forgetting. First, there is evidence that semantic and episodic retrieval competition are not so cleanly separable, at least in studies of inhibition. For instance, semantic retrieval practice has been shown to impair episodic representations of similar items (Bäuml, 2002), and part-set cuing of episodically presented items appears to impair access to semantically related competitors (Kimball & Bjork,

2002; see later section on recognition testing for further discussion). Second, to adopt episode-specific inhibition as a theory for all varieties of episodic retrieval-induced forgetting ignores a very plausible feature of retrieval: subjects can weight the different cues that they use flexibly, depending on the task. In some tasks, episodic context may be the most important cue to weight, whereas in others, the semantic category may be more diagnostic in guiding retrieval. If so, whether one observes episode-specific inhibition, or joint effects of inhibition on both episodic and semantic representations may hinge on the relative weighting of attention on contextual versus categorical cues. Finally, the relationship between episodic and semantic representations is at present not theoretically resolved: episodes may or may not be structurally distinguishable from their semantic counterparts. These issues remain to be explored in greater depth. Nevertheless, in any study looking at whether episodically induced inhibition may be observed on conceptual implicit memory tests, it would be prudent to entertain episode-specific inhibition as a theoretical mechanism that may contribute to performance.

The foregoing examples illustrate the central importance of considering the nature of the representation that is likely to be tapped by a particular variety of test, and how this representation may relate to the one likely to be subject to inhibition. Failure to find evidence of inhibition on a given test may not indicate a lack of inhibition in general; it may simply reflect a mismatch in the type of representation tapped by the test and that affected by inhibition. Nevertheless, although indirect tests may not be diagnostic of inhibition, such experiments do serve to define the scope of inhibitory effects induced by episodic retrieval practice, and the nature of the representations affected.

Masking through covert cuing effects. As described earlier, inhibition tends to generalize to novel test cues that are unrelated to the items receiving retrieval practice or to the practiced cues themselves—a property known as cue-independence. However, whether cue-independent forgetting will be observed may depend on whether subjects use covert cuing strategies to augment their recall on the final memory test. Consider, for example, a study by Anderson et al. (2000). In this study, subjects studied items such as Red-Blood and Red-Tomato, and later did retrieval practice on Red-Blood. On a delayed recall test, subjects were cued to recall Tomato with an extra-list category label and a letter stem (e.g., Food-T___) to see whether or not any inhibition that was induced by retrieval practice would generalize to the novel extralist test cue (see Anderson & Green, 2001; Johnson & Anderson, in press; Levy & Anderson, 2002; for other studies using extralist cuing). As predicted, significant inhibition was found, suggesting cue-independent impairment. However, when asked, on a post-experimental questionnaire, whether they tried to

augment their memory search by recalling earlier-studied categories, some subjects reported using this “covert cuing” strategy (the average rating was 2.68 on a 5 point scale). Thus, when given the extralist category cue Food T___, some subjects may have covertly recalled the category “Red Things,” and used these two categories jointly to recall items. Subjects who reported using this strategy showed modestly reduced inhibition effects, compared to subjects who did not use this strategy (a reduction of the inhibition effect by 3% in Experiment 1, and by 7% in Experiment 2). Given that the usefulness of covert cuing may have been limited by the timing constraints used in the test of that experiment (4 s per cue), these findings suggest that covert cuing may act to reduce inhibition under less constrained conditions.

The foregoing findings may be understood by considering the effects of cue priming discussed earlier. To the extent that practiced categories are made highly accessible by retrieval practice, subjects who engage in covert cuing are more likely to covertly generate the practiced categories than they are the baseline categories. As a result, when trying to recall inhibited items, subjects using this strategy should be more likely to have not one, but two category cues at their disposal, conferring a cuing advantage to those items, relative to baseline items. Thus, inhibition may be compensated by the differential availability of compound cuing. Such compensation would lead to an inaccurate measure of the amount of inhibition that had initially taken place (masking), and perhaps even undo that inhibition for the items retrieved by compound cues.

To reduce the likelihood of covert cuing contaminating recall performance in studies using the independent probe method, several strategies appear effective. First, subjects are less likely to use covert cuing when the extralist cues are, in general, strongly related to the target item; if most cues are poorly related, subjects may look for additional information to supplement their recall. Second, providing an item specific cue such as a letter stem focuses subjects on recalling a particular item. Third, limiting the amount of time that subjects have to recall the critical item discourages the use of complex search strategies such as covert cuing. Fourth, using a large number of studied categories makes it unlikely that subjects will be able to recall the relevant studied category, even if they try. Finally, administering post-experimental questionnaires to obtain subjective reports of covert cuing can help to assess whether the foregoing strategies were effective. In using the independent probe method to establish the theoretical property of cue independence, it is vital to consider how such strategies may affect performance.

Special issues in recognition testing. Initially, we believed that retrieval-induced forgetting would not occur on recognition memory tests (Anderson & Bjork, 1994). This expectation was based on analogies to other in-

hibitory phenomena such as directed forgetting and retroactive interference, which exhibit little impairment on recognition tests, and on the idea that presentation of the item itself would release it from its inhibited state (Anderson & Bjork, 1994). This perspective has proven to be mistaken. Significant retrieval induced forgetting has been found on recognition memory measures, both in the retrieval practice paradigm and in closely related procedures.

The first demonstration of retrieval-induced forgetting on a recognition test using the retrieval practice paradigm was reported by Anderson, De Kok, and Child (1997). Subjects participated in the standard retrieval practice procedure except that after the 20 min delay, subjects were given a yes/no recognition memory test for all of the exemplars they had studied instead of cued recall. In Experiment 1, subjects were tested with category–exemplar pairs, one at a time, and exemplars of a given category were tested in blocks of 12 (six targets and six highly similar distractors intermixed). As can be seen in Fig. 6, significant retrieval-induced forgetting was observed, regardless of whether the unpracticed competitors were tested before practiced items in their category (tested 1st) or after them (tested 2nd). Subsequent experiments provided evidence that this impairment also occurred when exemplars were presented without their category labels, and in randomized tests instead of tests using category blocks. Thus, retrieval-induced forgetting can be observed even when subjects are tested with the inhibited item presented intact, and do not have to generate the item from incomplete cues. Anderson et al. (1997) also observed within-category output interference on their recognition tests (to see this, compare tested first to tested second in Fig. 6 for each condition), consistent with other studies that have reported output interference on recognition tests (Smith, 1971). Similar inhibition effects have been

reported in two recent experiments by Hicks and Starns (in press) that used an item recognition test. (see Dopkins & Ngo, 2002, for a potentially related inhibition effect induced by incidental retrieval of an earlier presentation of an item during its repetition). Radvansky (1999) also found evidence for inhibition on a speeded recognition memory test using the fan effect procedure. In addition to generalizing these effects beyond categorical materials, Radvansky's study demonstrated that impairment is cue-independent, as predicted by the inhibition view.

Given the evidence for impairment on recognition tests, the question arises as to why such effects would occur for retrieval-induced forgetting and not other phenomena such as directed forgetting and retroactive interference. Although it is possible that retrieval-induced forgetting may be produced by qualitatively different mechanisms, other explanations exist. One possibility is that recognition tests might be most sensitive to retrieval-induced forgetting when the recognition judgments require active recollection rather than a mere assessment of familiarity. In the Anderson et al. study just discussed, subjects were asked to claim that they recognized an item only if they were very confident that it had occurred in the earlier study phase. These instructions should have encouraged a greater weight on recollective processes. If so, perhaps directed forgetting might also be found on recognition tests if tests required active recollection. Consistent with this idea, directed forgetting does cause impairment on recognition tests requiring subjects to make source memory judgments (e.g., Geiselman, Bjork, & Fishman, 1983). Thus, although simple yes/no recognition tasks appear to be insensitive to directed forgetting, judgments that require active retrieval of a particular episodic trace show the effect, as in recall paradigms. Although this account reconciles the patterns of inhibitory effects on recognition tests across the two paradigms, it leaves unexplained why familiarity-based judgments might fail to exhibit inhibition.

Another difficulty that may arise is the potential for the distractor items on a recognition test to be suppressed. For example, practicing Fruit Orange may suppress not only other studied items such as Fruit Banana, but also nonstudied items such as Fruit Strawberry. Because nonstudied exemplars are the very items that would be most useful to employ as distractors, both targets and distractor items may be impaired. Such effects ought to make it difficult to use signal detection methodology to measure inhibition. Consider the idealized familiarity distributions in Figs. 7A–C. Fig. 7A represents the situation before retrieval practice has taken place and shows familiarity distributions for baseline items and their distractors. Baseline items are assumed to be more familiar than distractors due to their recent presentation on the study list, and so the




Fig. 6. Retrieval-induced forgetting in recognition memory (Anderson et al., 1997). On a final category–exemplar pair recognition test, subjects were impaired in their ability to recognize unpracticed competitors, as measured by corrected recognition (hits–false alarms). This effect occurred regardless of whether unpracticed competitors were tested in the first half of their respective categories, or in the second half.

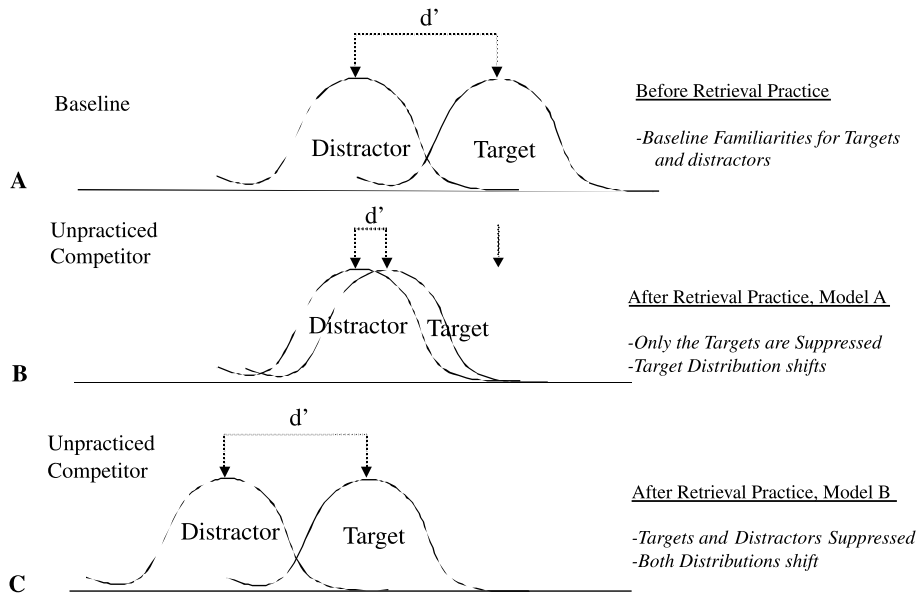


Fig. 7. An illustration of why it may sometimes be difficult to detect inhibition on recognition memory tests, in terms of signal detection theory. Each figure represents a continuum of familiarity values for items stored in memory, with distributions for studied targets and non-studied distractor items presented on the recognition test. (A) Before retrieval practice, all studied items are presumed to be more familiar to subjects than are distractors. (B) Familiarity distributions for unpracticed competitors and their distractors, after retrieval practice has been performed, according to Model A. In Model A, retrieval practice is presumed to selectively suppress the unpracticed competitors and NOT their corresponding, highly similar distractors in semantic memory. This leads target items to be less familiar, shifting the overall familiarity distribution for those items to the left, closer to the distribution for distractors, reducing d' . (C) The same familiarity distributions as plotted in (B), but plotted according to Model B. In Model B, retrieval practice is presumed to suppress both the unpracticed competitors and the highly similar distractors. This leads both target and distractor items to be less familiar, shifting the distributions for both to the left, leaving d' unaffected by suppression. Thus, if retrieval practice suppresses both unpracticed competitors and their distractors, impairment may not be observed on a recognition memory test, because d' will remain constant for baseline items and unpracticed competitors.

baseline distribution is shifted to the right. Figs. 7B and C represent the situation after retrieval practice, according to the view that retrieval practice: (1) suppresses only other episodically studied competitors and not distractors, or (2) suppresses both episodic and semantically related competitors that serve as distractors. If inhibition is restricted to episodically related competitors (7B), impairment should be measurable using d' because retrieval practice selectively shifts the distribution for unpracticed competitors, but not those of their distractors (note the leftward shift of the target distribution in Figs. 7B and C). No such shift occurs for baseline categories (Fig. 7A), so a difference in d' should emerge. However, if inhibition also affects unstudied semantically related competitors, both distributions will be shifted (Fig. 7C). Because d' only provides a measure of the relative discriminability of targets and distractors, inhibition may be quite difficult to measure, relative to baseline categories that have not shifted (see Samuel, 1996, for an analogous signal detection analysis in the context of speech perception; a similar point was also made in the context of the revelation effect by Hicks &

Marsh, 1998). Thus, inhibition may be difficult to detect on recognition tests not because inhibition has been released or does not affect familiarity, but because the nature of the test requires the use of foils that are themselves suppressed. Here again, the way in which the test is administered yields an inaccurate measure of how much inhibition truly took place, masking those effects.

There is good reason to suspect that inhibitory processes recruited during episodic retrieval suppress competing items in semantic memory. First, retrieval induced forgetting is a general phenomenon that occurs on both semantic and episodic retrieval tests (e.g., Blaxton & Neely, 1983; Johnson & Anderson, in press), showing that semantic representations are susceptible to inhibition. Second, inhibition effects have been previously shown to span episodic and semantic memory. Retrieving an exemplar of a category from semantic memory can suppress episodic memory for other exemplars that were studied previously (Bäuml, 2002). If semantic retrieval can suppress episodic memory, it seems likely that episodic retrieval might also suppress semantically related competitors that are not studied.

Finally, in a recent study, Kimball and Bjork (2002) found that presenting part-set cues during a recall test not only impaired remaining items that were studied in that set (as expected), but also reduced the intrusion rate for critical nonstudied semantic items that tend to be mistakenly recalled with those same materials. Taken together, these results suggest that the sensitivity of recognition tests to inhibitory effects may be masked by suppression of related semantic distractors.

Summary

The foregoing review highlights the core properties of retrieval-induced forgetting and some of its boundary conditions. Taken together, these properties argue for a strong parallel between selective retrieval and the more general situation of response override. In particular, the need to selectively retrieve a target item in the face of interference from one or more prepotent memories leads to the suppression of those memories, and this suppression underlies later forgetting of those items. Although inhibitory effects are sometimes moderated or masked by representational or testing factors, the basic finding is quite general and likely to underlie many cases of forgetting associated with interference. The experience of forgetting is more likely to be caused by inhibitory control processes that help to focus retrieval than by the strengthening of competing associations in memory.

Stopping retrieval through inhibitory control

In the preceding review, we discussed evidence for inhibitory processes in selective retrieval situations, which we argued are likely to require response override. However, response override is involved in other situations as well, such as when we need to stop a response from occurring at all. In memory retrieval, this ability could prove useful to prevent a particular memory from coming into consciousness. Indeed, we sometimes confront reminders of things that we would prefer not to think about: the sight of a car may remind us of an accident we had, or of a former significant other who drove that type of car; or the sight of the world trade center in an old movie may lead us to terminate the natural progression from cues to memories. Other times, we may wish to focus on a particular thought or idea without letting the mind wander. Can inhibitory control mechanisms be engaged to serve these goals? Can inhibition halt the retrieval process? If so, how? Anderson and Green (2001) recently looked at this issue by examining how stopping retrieval affected the memories that were to be retrieved. To study this, they developed a new procedure modeled after the widely used Go/No-Go task, which has been used to measure the ability to stop a prepotent motor response and to study its neural basis in both humans (e.g., Casey et al., 1997; de Zubicaray,

Andrew, Zelaya, Williams, & Dumanoir, 2000; Garavan, Ross, & Stein, 1999) and monkeys (e.g., Sakagami & Niki, 1994). In one version of the Go/No-Go task, letters are presented one at a time and subjects must press a button as quickly as possible whenever they see a letter, *except* when the letter is an X. When they see an X, they are supposed to avoid pressing the button. The majority of trials are designed to require a button press, so that when an X occurs, subjects have difficulty withholding their motor response. The ability to withhold the response is taken as a measure of inhibitory control.

To explore whether people can stop retrieval, Anderson and Green (2001) adapted the go/no-go task to create the *think/ no-think paradigm*. In this procedure, subjects studied pairs of weakly related words (e.g., flag—sword, ordeal—roach) and were then trained to provide the second word (e.g., roach; hereinafter referred to as the response word) whenever they were given the first word as a cue (e.g., ordeal). Subjects then entered the think/no-think phase, which required them to exert executive control over the retrieval process. For most of the trials in this phase, the task was the same as it had been during training—to recall and say aloud the corresponding word as quickly as possible at the sight of its retrieval cue. For certain cues, however, subjects were admonished to avoid thinking of the response word. It was emphasized that it was not enough to avoid saying the response word—it was crucial on those trials to prevent the associated memory from entering *conscious awareness* at all. Thus, subjects had to override not only a vocal motor response, but also the cognitive act of retrieval. Could subjects recruit inhibitory control mechanisms to stop the memory from entering consciousness?

Of course, Anderson and Green could not directly measure whether subjects stopped the memory from entering consciousness, but if inhibitory mechanisms were recruited, later recall of the excluded memory should be impaired. To examine this, immediately after the think/no-think phase, subjects were given the cues for all of the pairs, but they were now asked to recall the response for each of them. As expected, forgetting occurred: response words that subjects tried to keep out of awareness were impaired compared to baseline pairs they had studied initially but had not seen during the think/ no-think phase. The more often subjects tried to stop retrieval, the worse recall for the excluded memory became (see Fig. 8A). Interestingly, avoided words were harder to recall even though subjects had encountered as many as 16 reminders (i.e., cues) during the think/ no-think phase. Under normal circumstances, reminders would be expected to facilitate the reminded memory, much as it did for the items to which subjects continued to respond (Fig. 8A). Anderson and Green (2001) further established that this impairment was cue

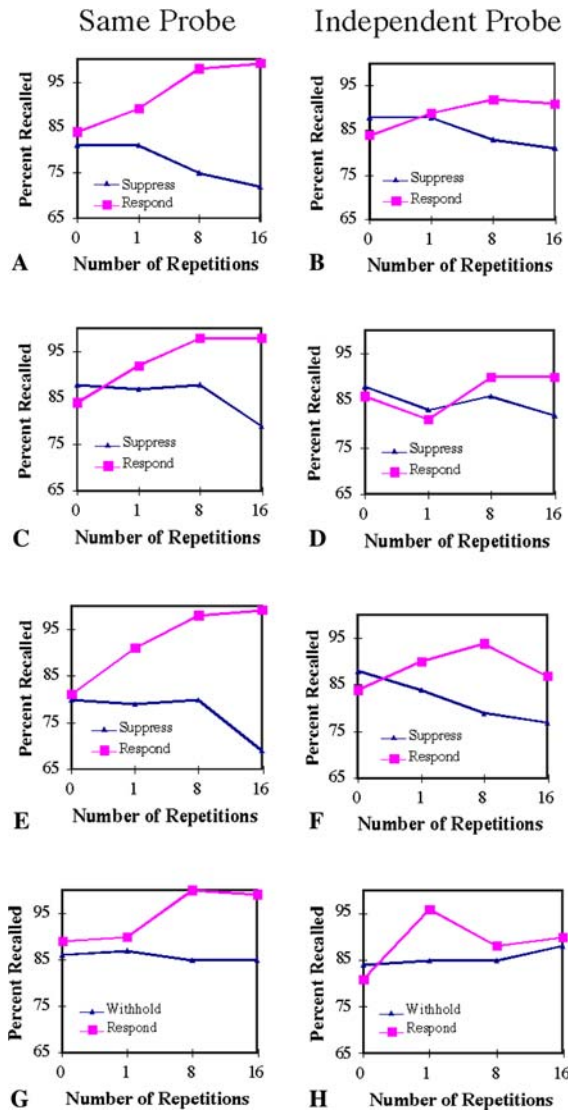


Fig. 8. Final recall performance in four experiments reported by Anderson and Green (2001) using the think/no-think procedure. Each plot represents the percentage of items that subjects recalled on the final recall test as a function of the number of times that they suppressed the item (suppression condition), or tried to recall it (respond). The left panel in each row represents final recall performance when tested with the originally trained retrieval cue (i.e., the “Same probe”), whereas the right panel in each row represents final recall performance when tested with a novel, independent, extralist category cue. (A and B) depicts performance in Experiment 1; (C and D) depicts performance in an experiment offering monetary incentives, and encouraging guessing on the final test; (E and F) depicts performance when subjects were misled regarding the expected outcome of the study just before the test; (G) depicts final test performance when subjects are simply asked, during the think/no-think phase to simply not say the response word (withhold) instead of to not think about it; final memory is not impaired.

independent, echoing the results of Anderson and Spellman (1995): forgetting occurred regardless of whether subjects were tested with the originally studied cue word (e.g., ordeal) or with a novel independent cue never studied in the experiment (e.g., insect r___ for roach; Fig. 8B). This cue-independence argues that the forgetting is not caused solely by associative interference; rather, impairment reflects active suppression of the excluded memory itself.

Anderson and Green (2001) ruled out the possibility that subjects might have deliberately withheld answers on the final test due to confusion or to expectations about the purpose of the experiment. In one experiment, subjects were told that they would be paid for all correct answers and were urged to respond to every cue, even if they were guessing. Another group was misled to believe that the experimenters expected that their memory would be better for words they had avoided thinking about. Both manipulations left the inhibition pattern unchanged (see Figs. 8C–D for recall performance in the original cue and independent cue conditions respectively for the monetary incentives experiment; see Figs. 8E and F, for the same conditions for performance by misled subjects), demonstrating that subjects were neither confused nor purposefully withholding responses. In a final experiment, subjects were merely asked to avoid saying the response out loud and all mention of preventing it from entering awareness was eliminated. No inhibition was observed (Fig. 8G), indicating that the recall deficits in the preceding experiments were not merely due to suppression of the vocal response for avoided words. These results isolate forgetting in the think/no-think paradigm to processes directed at keeping the unwanted declarative memory out of awareness and demonstrate that this cognitive act has persisting consequences for the avoided memories.

The impaired memory observed by Anderson and Green (2001) suggests that inhibitory control mechanisms may be recruited when we seek to regulate awareness of unpleasant or intrusive memories. In particular, whenever the environment is such that it presents unavoidable reminders to something that we would prefer not to think about, people may resort to controlling their memories instead. The end result may be impaired memory for the things that people avoid thinking about. This suggests that the think/no-think paradigm of Anderson and Green (2001) may provide a useful laboratory model of the voluntary form of repression (suppression) proposed by Freud (Freud, 1966). If so, results from this paradigm and other related paradigms such as the directed forgetting procedure may have implications for understanding clinical phenomenon relating to motivated forgetting (Anderson, 2001; Anderson & Green, 2001; Bjork, Bjork, & Anderson, 1998; Conway, Harries, Noyes, Racsmany, & Frankish, 2000; Deprince & Freyd, 2001; Myers, Brewin, & Power,

1998; see Golding & MacCleod, 1998 for a review of directed forgetting).

Relationship to classical interference theories of forgetting

Although the executive control view is a relatively new approach to interference, many of its aspects resemble components of classical interference theory. In this section, I discuss some of the specific relations between this view and four mechanisms discussed in classical interference theory: response competition, unlearning, reciprocal inhibition, and response-set suppression. The executive control approach validates many of the intuitions behind these classical proposals, while at the same time questioning the historical emphasis that has been placed on associative learning as a source of forgetting.

McGeoch's response competition theory

According to McGeoch's classical response competition theory, attaching more than one response to a retrieval cue leads those responses to compete with one another when the cue is presented later on. The more competing responses, or the stronger a competing response becomes, the more difficult it should be to recall a given item. McGeoch's emphasis on the importance of sharing a retrieval cue as a condition of interference was inherited from Müller and Pilzecker (1900), and continues today in the form of relative strength or ratio-rule models of retrieval (e.g., Anderson, 1983; Mensink & Raaijmakers, 1988). In essence, these theories posit that the addition of new structure into memory leads to the occlusion or blocking of a target event.

Many of the basic assumptions of McGeoch's response competition theory are accepted in the current executive control approach. For instance, the presentation of a retrieval cue is presumed to activate all associated responses according to their strengths of association to the cue and these responses are thought to compete with one another for access to conscious awareness. It is this retrieval competition that precipitates the need for executive control. According to the executive control approach, however, this competition is usually not enough by itself to impair memory recall for a target because inhibitory processes may be deployed to overcome the competition. Furthermore, the empirical relationship between the number of competing responses and the probability of recalling a target item is also accepted by the theory, along with the notion that strengthening a competing response is empirically associated with a decrement in recall for a target.

Where the executive control approach advanced here differs from McGeoch's theory is in the underlying mechanism that produces these relationships. According

to the executive control approach, the probability of recalling a target item does not automatically decrease as a consequence of adding new associations, or as a consequence of strengthening a competing association. Structural changes may impair the later recall of a target item if they increase the chances that nontarget items will occasionally be retrieved before the critical target. To the extent that competitors are retrieved earlier, the target will be suppressed at output. The probability that this suppression will impair target performance should go up with the number of competitors because this will lead more competitors to be recalled before the target, on average. By this view then, strengthening a competitor should not impair target recall provided that the target can be ensured to be tested before the competitor, a finding that has been observed many times now in experiments evaluating the hypothesis of strength-dependent forgetting. Thus, it is not the addition of new associations, nor their strengthening that impairs memory, but rather the increased likelihood of suppression correlated with those structural changes.

There are several circumstances, however, in which response competition might impair memory. First, whenever a cue is presented that is associated to a stronger and a weaker response and the subject is told to only report the first thing that comes to mind, response competition might underlie interference effects. Naturally, if the subject is to report the first thing that comes to mind, the stronger response will typically prevail over the weaker one, causing the omission of the latter. This will lend the appearance of inaccessibility of the weaker response when it may not be inaccessible at all. Second, when the subject is given a very short time to make memory responses to a cue, interference may be produced by blocking. Here again, stronger responses will leap to mind most readily and potentially use up all the time that the subject has to express their knowledge of the associated memories. Even if all responses are of equal strength, the addition of new responses might increase the chances that some nontarget item will be reported to the exclusion of a target in a limited time window. In both of these cases, interference effects may reflect some combination of suppression arising from the prior output of nontarget items and blocking produced by insufficient time to express available knowledge. In fact, much of the early work on interference theory up until the late 1950s employed the modified free-recall test (i.e., the MFR test), which required that the subject provide only a single response in a limited time window. With the advent of the modified-modified-free recall test (MMFR), subjects were asked to recall all available responses and were given a longer period to recall them (Barnes & Underwood, 1959), a procedure that was thought to provide a better test of the true availability of responses in memory. Third, when the measure of interference is reaction time, the presence of multiple

competitors or a single strong competitor should slow the recall of a target; again, this retrieval interference is thought to be an essential step in triggering inhibitory control. Finally, special populations with deficits in executive function (e.g., older adults, children, frontal-lobe damaged patients) may be sufficiently challenged in inhibiting competitors so that retrieval competition cannot be effectively managed. Under these circumstances, competitors may block the retrieval of a target and cause impairment that does not reflect inhibitory processes.

Melton and Irwin's unlearning theory

According to the unlearning hypothesis, interference effects are caused in part by the unlearning of associative connections linking a retrieval cue to a response. Specifically, when a person is trying to recall a newly learned response (e.g., the new phone number for a friend), previously learned responses to that same cue (e.g., the old phone number for that friend) may sometimes be elicited accidentally. Elicitation could take the form of an overt or covert intrusion of the unwanted item. To the extent that the older response is incorrect, it was thought to go “unreinforced,” and therefore suffer extinction effects analogous to those exhibited by animals in conditioning experiments. Associative unlearning was a critical component of Melton and Irwin's classical two-factor theory of interference (Melton & Irwin, 1940), which also incorporated response competition. The modern descendants of this view include the many connectionist learning systems that might attribute forgetting in part to the alteration of weights between representational units.

The current approach shares much with the unlearning hypothesis: it focuses on the intrusion of unwanted memory responses during retrieval as a condition leading to the forgetting of the intruding items; it posits a process that responds to intrusions in such a way as to render them less likely in the future—changing some aspect of the intrusion's representation. Thus, a special forgetting process is proposed. It differs, however, both in its theoretical orientation, and in the nature of the forgetting mechanism. The unlearning idea was a theoretical analogy inspired by the behaviorist learning approach. Simple, automatic processes were proposed: learning was the positive adjustment of associations, forgetting, the negative adjustment. The executive control approach, however, is concerned with the moment-by-moment control of behavior with respect to flexible goals. It assumes mechanisms by which mental representations are adjusted dynamically in contexts in which their ongoing accessibility might disrupt our aims. The mechanisms that achieve this adjustment are not thought of as general learning processes, but as processes that control the operational state of a system.

These different orientations lead to different conceptualizations of how intruding memories become im-

paired: whereas unlearning posits a decrement in the associative bond linking a cue to a target, the executive control approach attributes impairment to a suppression of the target itself. Thus, the current approach predicts cue-independent impairment, whereas unlearning does not. The existence of cue-independent impairment of course does not rule out the possibility that associative unlearning might also occur and contribute to the impairment observed in both retrieval-induced forgetting and classical interference paradigms.

Osgood's reciprocal inhibition hypothesis

An often overlooked theory is Osgood's reciprocal inhibition approach to interference. According to this theory, strengthening the association between a stimulus and a response also strengthens an inhibitory association between the stimulus and semantically antagonistic responses that are attached to it (Osgood, 1946, 1948). For instance, if subjects learn the pair *Tree-Elated*, a positive association is formed between the two words, but an inhibitory one is also established between *Tree* and the antagonistic response *Dejected*. In essence, subjects not only learn to make the correct response, but also to NOT make the opposite response—a notion borrowed from Hull's behavioral theory (Hull, 1943). Both the excitatory and inhibitory associations were thought to generalize semantically, so that intermediate responses such as *Low*, would also suffer inhibition, by virtue of its similarity to *Dejected*. Osgood provided some support for this theory, showing gradually increasing retroactive interference across similar, neutral, and antagonistic responses to stimuli, as a result of interpolated associative learning.

Osgood's theory is perhaps the first theory of retroactive interference that attributed impairment to an inhibitory mechanism. In Osgood's framework, impairment was thought to be a direct result of inhibiting the potentially intrusive response, and so the theory can explain cue independent impairment. Here again, the hypothesis bears some resemblance to the current executive control theory. However, Osgood's assertion that inhibition is a direct function of semantic antagonism between two responses is not a feature of the current approach, nor is there any commitment to the development of an inhibitory association between a stimulus and an unwanted response, as Osgood proposed. In the current perspective, if a cue activates a memory that is unwanted—either because it interferes with a retrieval attempt, or because it is distracting or unpleasant—inhibitory control mechanisms can be recruited to suppress the item. Consistent attempts to suppress a memory may or may not result in the formation of an “inhibitory habit” for a given item, as Osgood proposes, but this possibility is beyond the scope of the present theory.

Postman's response-set suppression hypothesis

Near the end of the classical interference era, Postman and colleagues (Postman et al., 1968) proposed a theory of interference that departed substantially from approaches previously proposed. As highlighted in the preceding sections, most classical accounts of interference were embedded within larger scale theories of associative learning that had their conceptual roots in behaviorist learning theory. Forgetting was assumed to reflect the effects of competition between alternate responses, or the degradation of associations by general learning mechanisms. However, Postman proposed mechanisms that went well beyond the somewhat limited conceptual arsenal of most learning frameworks. According to his response-set suppression hypothesis, retroactive interference was caused by the active suppression of response members from the initial list. Suppression was thought to occur during the acquisition of the second list of pairs by what Postman referred to as a "selector mechanism." The function of the selector mechanism was to both enhance the representations of responses that were intended to be part of the current response set and to suppress outdated response sets. The suppression process helped to reduce proactive interference caused by the initial list, and to effectively "shift" into a "response set" more appropriate to the current task.

The response-set suppression hypothesis can be seen an early example of the executive control approach. Like the executive control theory, this hypothesis attributed forgetting to a mechanism that directly suppressed the response representations of items from the first list of pairs. This mechanism was clearly linked to response override: it helped the organism to "select" current, more contextually appropriate response sets in the face of interference from preceding response sets. Thus, this hypothesis acknowledged the need to control memory in accordance with current goals, and advocated a special process to achieve that control. Nevertheless, the current hypothesis differs from Postman's theory in several respects. First, according to the response-set suppression view, the selector mechanism was thought to act on entire "response repertoires" and not at the level of individual responses. So, if a subject learned a list of ten pairs, followed by a second list of ten pairs, all responses from the initial list would be suppressed, irrespective of whether or not the stimulus member for a given first-list item was also used in the second list. The set of first list responses was suppressed as a whole, and the set of second list responses, facilitated. The current approach is more flexible, permitting for suppression of specific competing responses. Accordingly, it should be (and is) possible to suppress only select items from a list, based on how much interference they cause during retrieval of second list items, as is evident in studies of retrieval-induced forgetting. Second, the response-set suppression

view drew a tight connection between the need to facilitate a new response set in order to suppress a preceding set. The current approach entertains the idea that suppression can be directly applied to an unwanted memory without the need to facilitate a competing response or response set. Work with the think/no-think paradigm, for example, suggests that suppression is directly applied to distracting memories. However, further work needs to be done to determine whether such direct suppression is truly possible. Finally, Postman's theory made a variety of specific assumptions intended to explain the conditions under which spontaneous recovery from retroactive interference should occur. Although these assumptions may be correct, they are not an intrinsic part of the current theory, as it is presently specified.

Despite these differences, the present theory might be regarded as a modern cousin to response-set suppression that decouples it from the particular paradigm within which the theory was developed. The response-set suppression view has been overlooked as an approach to interference in part because the theory was developed towards the end of the classical interference era, when the field became captivated by cognitive theory. The shift away from interference research led to the abandonment of the theory, and of research on interference generally. Ironically, to the extent that interference was discussed after the cognitive revolution, theories became far more associationistic than Postman's—more in the vein of McGeoch's response competition theory (e.g., Anderson, 1983; Mensink & Raajimakers, 1988; Rundus, 1973). The developing interest in executive control functions in the last 15 years, and work on inhibitory processes has made it possible to view Postman's theory in a different light.

Concluding remarks

Research on interference has occupied a central role in the science of memory since the beginnings of experimental psychology. Throughout much of this long history, theoretical discussions of interference have been dominated by ideas that were either directly borrowed from, or inspired by classical associative learning theories. In many ways, this conceptual influence pervades thinking about interference even today not only in how this phenomenon is described in modern textbooks, but also in how it is explained within current theoretical frameworks. In modern textbooks, retroactive interference is often defined, for example, as the forgetting that arises as a result of new *learning*, and proactive interference, as the forgetting that arises as a result of previous *learning*. If a theory is described at all, it is often the classical two-factor theory of Melton and Irwin (1940). Despite many differences in terminology and constructs particular to cognitive psychology, current theoretical accounts of interference have essentially

returned to McGeoch's associative interference theory. There are excellent reasons for the continuing influence of these classical ideas about learning: interference effects are highly correlated with the storage of new traces into memory, and with the modification of existing ones. The act of learning a new list does impair memory for a previous one, and strengthening a competing association is often associated with impaired recall of related traces. These empirical relationships lend force to the idea that forgetting ultimately derives from the ever changing contents of memory, and our inability to cope with the competition created by those changes.

In this article, I have argued that despite these empirical relationships, we should rethink our view of how interference leads to forgetting. I have argued that a theory of interference should be framed in the larger context of how organisms control their own thoughts and actions. Memory retrieval is just a special case of a broad class of situations that recruit executive control processes, and it is these processes—particularly inhibition—that cause forgetting. By this view, the empirical relationship between associative learning and forgetting that emerged with Müller and Pilzecker (1900) and that drives theorizing today, should not be construed as proof that new learning impairs memory, as posited in many classical and modern models. New learning sets the stage for the mechanism that actually causes forgetting: inhibition. Inhibition is triggered as a direct response to the competition caused by related traces and the goal to selectively retrieve a target, or—in the case of motivated forgetting—in response to the goal to prevent awareness of a distracting memory. The forgetting that results is not a passive side effect of the new learning, but a consequence of the mechanisms that have evolved to allow organisms to override prepotent responses. These mechanisms are essential to our ability to behave in a flexible, context-appropriate manner.

I have argued that this view validates many of the insights offered by classical theories, while questioning the widespread assumption that forgetting is tied in a direct way to new associative learning. Rather, our experiences of forgetting—of past experiences, of our friends' names, or of ideas with which we were once adept, are seen as costs of the very mechanisms that enable us to direct cognition to internal thoughts and to the external environment.

Acknowledgments

I would like to thank Douglas Hintzman, Larry Jacoby, Ben Levy, Brice Kuhl, Sarah Johnson, and two anonymous reviewers for comments on an earlier version of the manuscript. Portions of this article evolved from a planned paper to be entitled "How to not find inhibition in memory without really trying" by

the present author and B.A. Spellman. Preparation of this paper was supported by a grant from the National Institute of Mental Health.

References

- Allen, G. A., Mahler, W. A., & Estes, W. K. (1969). Effects of recall tests on long-term retention of paired associates. *Journal of Verbal Learning and Verbal Behavior*, 8, 463–470.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R., & Reder, L. M. (1999). The fan effect: New results and new theories. *Journal of Experimental Psychology: General*, 128, 186–197.
- Anderson, M. C. (1989). *Similarity-based retrieval inhibition and mnemonic techniques*. Unpublished Master's Thesis, University of California, Los Angeles, Los Angeles.
- Anderson, M. C. (2001). Active forgetting: Evidence for functional inhibition as a source of memory failure. *Journal of Aggression, Maltreatment, and Trauma*, 4, 185–210.
- Anderson, M. C., & Bell, T. (2001). Forgetting our facts: The role of inhibitory processes in the loss of propositional knowledge. *Journal of Experimental Psychology, General*, 130, 544–570.
- Anderson, M. C., Bjork, E. L., & Bjork, R. A. (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., & Bjork, R. A. (1994). Mechanisms of inhibition in long-term memory: A new taxonomy. In D. Dagenbach, & T. Carr (Eds.), *Inhibitory processes in attention, memory and language* (pp. 265–326). San Diego: Academic Press.
- Anderson, M. C., De Kok, D., & Child, C. (1997). Retrieval-induced forgetting on a test of recognition memory [Abstract]. *Abstracts of the Psychonomic Society*, 2, 2.
- Anderson, M. C., & Green, C. (2001). Suppressing unwanted memories by executive control. *Nature*, 410, 131–134.
- Anderson, M. C., Green, C., & McCulloch, K. C. (2000). Similarity and inhibition in long-term memory: Evidence for a two-factor model. *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.), *Memory. Handbook of perception and cognition* (2nd ed., pp. 237–313). San Diego, CA: Academic Press.
- Anderson, M. C., Ochsner, K., Kuhl, B., Cooper, J., Robertson, E., Gabrieli, S. W., Glover, G., & Gabrieli, J. D. E. (submitted). Neural systems underlying the suppression of unwanted memories.
- Anderson, M. C., & Shivde, G. S. (in preparation a). Inhibition in episodic memory: Evidence for a retrieval-specific mechanism.

- Anderson, M. C., & Shivde, G. S. (in preparation b). Strength is not enough: Evidence against a blocking theory of retrieval-induced forgetting.
- 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.
- Arbuckle, T. Y. (1966). Intratrial interference in “immediate” memory. *Dissertation Abstracts*, *26*, 5554.
- Balota, D. A. (1994). Visual word recognition: The journey from features to meaning. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics* (pp. 303–358). San Diego, CA, US: Academic Press.
- Barnes, J. M., & Underwood, B. J. (1959). “Fate” of first-list associations in transfer theory. *Journal of Experimental Psychology*, *58*, 97–105.
- Bäuml, K. (1996). Revisiting an old issue: Retroactive interference as a function of the degree of original and interpolated learning. *Psychonomic Bulletin & Review*, *3*, 380–384.
- Bäuml, K. (1997). The list-strength effect: Strength-dependent competition or suppression. *Psychonomic Bulletin & Review*, *4*, 260–264.
- Bäuml, K. (1998). Strong items get suppressed, weak items do not: The role of item strength in output interference. *Psychonomic Bulletin & Review*, *5*, 459–463.
- Bäuml, K. (2002). Semantic recall can cause episodic forgetting. *Psychological Science*, *13*, 356–360.
- Bäuml, K., & Hartinger, A. (2002). On the role of item similarity in retrieval-induced forgetting. *Memory*, *10*, 215–224.
- Barnhardt, T. M., Glisky, E. L., Polster, M. R., & Elam, L. (1996). Inhibition of associates and activation of synonyms in the rare-word paradigm: Further evidence for a center-surround mechanism. *Memory & Cognition*, *24*, 60–69.
- Bjork, E. L., Bjork, R. A., & Anderson, M. C. (1998). Varieties of goal-directed forgetting. In J. M. Golding, & C. M. MacLeod (Eds.), *Intentional forgetting: Interdisciplinary approaches* (pp. 103–137). Mahwah, NJ: Erlbaum.
- Bjork, R. A. (1975). Retrieval as a memory modifier. In R. Solso (Ed.), *Information processing and cognition: The Loyola symposium* (pp. 123–144). Hillsdale, NJ: Erlbaum.
- Bjork, R. A. (1989). Retrieval inhibition as an adaptive mechanism in human memory. In H. L. Roediger III, & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 309–330). Hillsdale, NJ: Erlbaum.
- Blaxton, T. A., & Neely, J. H. (1983). Inhibition from semantically related primes: Evidence of a category-specific retrieval inhibition. *Memory & Cognition*, *11*, 500–510.
- Butler, K. M., Williams, C. C., Zacks, R. T., & Maki, R. H. (2001). A limit on retrieval-induced forgetting. *Memory & Cognition*, *27*, 1314–1319.
- Carrier, M., & Pashler, H. (1992). The influence of retrieval on retention. *Memory & Cognition*, *20*, 633–642.
- Casey, B. J., Trainor, R. J., Orendi, J. L., Schubert, A. B., Nystrom, L. E., Giedd, J. N., Castellanos, F. X., Haxby, J. V., Noll, D. C., Cohen, J. D., Forman, S. D., Dahl, R. E., & Rapoport, J. L. (1997). A developmental functional MRI study of prefrontal activation during performance of a Go-No-Go task. *Journal of Cognitive Neuroscience*, *9*, 835–847.
- Ciranni, M. A., & Shimamura, A. P. (1999). Retrieval-induced forgetting in episodic memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 1403–1414.
- Conway, A. R., & Engle, R. W. (1994). Working memory and retrieval: a resource-dependent inhibition model. *Journal of Experimental Psychology: General*, *123*, 354–373.
- Conway, M. A., Harries, K., Noyes, J., Racsmany, M., & Frankish, C. R. (2000). The disruption and dissolution of directed forgetting: Inhibitory control of memory. *Journal of Memory and Language*, *43*, 409–430.
- Dagenbach, D., Carr, T. H., & Barnhardt, T. M. (1990). Inhibitory semantic priming of lexical decisions due to failure to retrieve weakly activated codes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 328–340.
- Deprince, A. P., & Freyd, J. J. (2001). Memory and dissociative tendencies: The roles of attentional context and word meaning in a directed forgetting task. *Journal of Trauma and Dissociation*, *2*(2), 67–82.
- de Zubicaray, G. I., Andrew, C., Zelaya, F. O., Williams, S. C. R., & Dumanoir, C. (2000). Motor response suppression and the prepotent tendency to respond: A parametric fMRI study. *Neuropsychologia*, *38*, 1280–1291.
- Dodson, C. S., Johnson, M. K., & Schooler, J. W. (1997). The verbal overshadowing effect: Why descriptions impair face recognition. *Memory & Cognition*, *25*, 129–139.
- Dong, T. (1972). Cued partial recall of categorized words. *Journal of Verbal Learning and Verbal Behavior*, *11*, 654–661.
- Dopkins, S., & Ngo, C. T. (2002). Inhibition of verbal memory retrieval as a consequence of prior retrieval. *Journal of Memory and Language*, *46*, 606–621.
- 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.
- Finger, K., & Pezdek, K. (1999). The effect of cognitive interview on face identification accuracy: Release from verbal overshadowing. *Journal of Applied Psychology*, *84*, 340–348.
- Fisher, R. P., & Craik, F. I. M. (1977). Interaction between encoding and retrieval operations in cued recall. *Journal of Experimental Psychology: Human Learning and Memory*, *3*, 701–711.
- Freud, S. (1966). In J. Strachey (Ed.), *The standard edition of the complete psychological works of Sigmund Freud* (pp. 117–128). Hogarth.
- Garavan, H., Ross, T. J., & Stein, E. A. (1999). Right hemispheric dominance of inhibitory control: An event-related functional MRI study. *Proceedings of the National Academy of Sciences USA*, *96*, 8301–8306.
- Gardiner, J. M., Craik, F. I., & Bleasdale, F. A. (1973). Retrieval difficulty and subsequent recall. *Memory & Cognition*, *1*, 213–216.
- Geiselman, R. E., Bjork, R. A., & Fishman, D. L. (1983). Disrupted retrieval in directed forgetting: A link. *Journal of Experimental Psychology: General*, *112*, 58–72.
- Gernsbacher, M. A., & Faust, M. E. (1991). The mechanism of suppression: A component of general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 245–262.
- Golding, J. M., & MacCleod, C. M. (Eds.). (1998). *Intentional forgetting: Interdisciplinary approaches*. Mahwah, NJ: Erlbaum.

- Hayes-Roth, B. (1977). Evolution of cognitive structures and processes. *Psychological Review*, 84, 260–278.
- Hicks, J. L., & Starns, J. (in press). Retrieval-induced forgetting occurs in tests of item recognition. *Psychonomic Bulletin & Review*.
- Hicks, J. L., & Marsh, R. L. (1998). A decrement-to-familiarity interpretation of the revelation effect from forced-choice tests of recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 1105–1120.
- Horton, D. L., & Kjeldgaard, P. M. (1961). An experimental analysis of associative factors in mediated generalizations. *Psychological Monographs*, 75, 26.
- Hull, C. L. (1943). *Principles of behavior: An introduction to behavior theory*. New York: Appleton-Century.
- Hunt, R. R., & Smith, R. E. (1998, November). *Retrieval may or may not cause forgetting*. Paper presented at the annual meeting of the Psychonomic Society, Dallas, TX.
- Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General*, 110, 306–340.
- Jenkins, J. J. (1963). Mediated associations: Paradigms and situations. In C. N. Cofer, & B. S. Musgrave (Eds.), *Verbal behavior and learning: Problems and processes* (pp. 210–257). New York: McGraw-Hill.
- Johnson, S., & Anderson, M. C. (in press). The role of inhibitory control in forgetting semantic knowledge. *Psychological Science*.
- Kimball, D. R., & Bjork, R. A. (2002). Influences of intentional and unintentional forgetting on false memories. *Journal of Experimental Psychology: General*, 131, 116–130.
- Kjeldgaard, P. M. (1968). Transfer and mediation in verbal learning. In T. R. Dixon, & D. L. Horton (Eds.), *Verbal behavior and general behavior theory* (pp. 1071–1075). Englewood Cliffs, NJ: Prentice Hall.
- Koutstaal, W., Schacter, D. L., Johnson, M. K., & Galluccio, L. (1999). Facilitation and impairment of event memory produced by photograph review. *Memory & Cognition*, 27, 478–493.
- Levy, B., Reinholz, J., & Anderson, M. C. (in preparation). Direct cuing of inhibitory processes in the Think/no-Think procedure.
- Levy, B. J., & Anderson, M. C. (2002). Inhibitory processes and the control of memory retrieval. *Trends in Cognitive Science*, 6, 299–305.
- Lustig, C., Hasher, L., & Toney, S. T. (2001). Inhibitory control over the present and the past. *European Journal of Cognitive Psychology*, 13, 107–122.
- MacLeod, M. (2002). Retrieval-induced forgetting in eyewitness memory: Forgetting as a consequence of remembering. *Applied Cognitive Psychology*, 16, 135–149.
- MacLeod, M. D., & Macrae, C. N. (2001). Gone but not forgotten: The transient nature of retrieval-induced forgetting. *Psychological Science*, 12, 148–152.
- Macrae, C. N., & MacLeod, M. D. (1999). On recollections lost: When practice makes imperfect. *Journal of Personality and 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.
- McDaniel, M. A., Friedman, A., & Bourne, L. E. (1978). Remembering the levels of information in words. *Memory & Cognition*, 6, 156–164.
- McGeoch, J. A. (1942). *The psychology of human learning: An introduction*. New York: Longmans.
- McGeoch, J. A., & Irion, A. L. (1952). *The psychology of human learning* (2nd ed.). New York: Longmans, Green, & Co.
- Meissner, C. A., & Brigham, J. C. (2001). The influence of retrieval processes in verbal overshadowing. *Memory & Cognition*, 29, 176–186.
- Melcher, J. M., & Schooler, J. W. (1996). The misremembrance of wines past: Verbal and perceptual expertise differentially mediate verbal overshadowing of taste memory. *Journal of Memory and Language*, 35, 231–245.
- Melton, A. W., & Irwin, J. M. (1940). The influence of degree of interpolated learning on retroactive inhibition and the overt transfer of specific responses. *American Journal of Psychology*, 3, 173–203.
- Mensink, G. J. M., & Raaijmakers, J. G. W. (1988). A model of interference and forgetting. *Psychological Review*, 95, 434–455.
- Miyamoto, A., & Anderson, M. C. (in preparation). Suppressing awareness: Evidence for a link between the regulation of consciousness and inhibitory phenomena in memory. *Psychological Science*.
- Moulin, C. J. A., Perfect, T. J., Conway, M. A., North, A. S., Jones, R. W., & James, N. (2002). Retrieval-induced forgetting in Alzheimer's disease. *Neuropsychologia*, 40, 862–867.
- Morris, C. D., Bransford, J. P., & Franks, J. J. (1977). Levels of processing versus transfer appropriate processing. *Journal of Verbal Learning and Verbal Behavior*, 16, 519–533.
- Müller, G. E., & Pilzecker, A. (1900). Experimentelle Beiträge zur Lehre vom Gedächtnis. *Zeitschrift für Psychologie*, 1, 1–300.
- Myers, L. B., Brewin, C. R., & Power, M. J. (1998). Repressive coping and the directed forgetting of emotional material. *Journal of Abnormal Psychology*, 107, 141–148.
- Nader, A., Coles, M. E., Brigidi, B., & Foa, E. B. (2001). The effect of practice on recall of emotional information in individuals with generalized social phobia. *Journal of Abnormal Psychology*, 110, 76–82.
- Nickerson, R. S. (1984). Retrieval inhibition from part-set cuing: A persisting enigma in memory research. *Memory & Cognition*, 12, 531–552.
- Norman, K. A., & O'Reilly, R. C. (in press). Modeling hippocampal and neocortical contributions to recognition memory: A complementary learning systems approach. *Psychological Review*.
- Osgood, C. E. (1946). Meaningful similarity and interference in learning. *Journal of Experimental Psychology*, 36, 277–301.
- Osgood, C. E. (1948). An investigation into the causes of retroactive interference. *Journal of Experimental Psychology*, 38, 132–154.
- Osgood, C. E. (1949). The similarity paradox in human learning: A resolution. *Psychological Review*, 56, 132–143.
- Perfect, T. J., Moulin, C. J. A., Conway, M. A., & Perry, E. (2002). Assessing the inhibitory account of retrieval-induced forgetting with implicit-memory tests. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 1111–1119.

- Postman, L. (1971). Transfer, interference, and forgetting. In J. W. Kling, & L. A. Riggs (Eds.), *Woodworth and Schlosberg's: Experimental psychology* (3rd ed., pp. 1019–1132). New York: Holt, Rinehart & Winston.
- Postman, L., Stark, K., & Fraser, J. (1968). Temporal changes in interference. *Journal of Verbal Learning and Verbal Behavior*, 7, 672–694.
- Raaijmakers, J. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, 88, 93–134.
- Rajaram, S., Srinivas, K., & Travers, S. (2001). The effects of attention on perceptual implicit memory. *Memory & Cognition*, 29, 920–930.
- Radvansky, G. A. (1999). Memory retrieval and suppression: The inhibition of situation models. *Journal of Experimental Psychology: General*, 128, 563–579.
- Radvansky, G. A., & Zacks, R. T. (1991). Mental models and the fan effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 940–953.
- Roediger III, H. L. (1973). Inhibition in recall from cueing with recall targets. *Journal of Verbal Learning and Verbal Behavior*, 12, 644–657.
- Roediger III, H. L., & McDermott, K. B. (1993). Implicit memory in normal human subjects. In F. Boller, & J. Grafman (Eds.), *Handbook of Neuropsychology* (Vol. 8, pp. 63–131).
- Roediger III, H. L., & Neely, J. H. (1982). Retrieval blocks in episodic and semantic memory. *Canadian Journal of Psychology*, 36, 213–242.
- Roediger, H. L., & Schmidt, S. R. (1980). Output interference in the recall of categorized and paired associate lists. *Journal of Experimental Psychology: Human Learning and Memory*, 6, 91–105.
- Roediger, H. L., Weldon, M. S., Stadler, M. A., & Rieger, G. H. (1992). Direct comparison of word stems and word fragments in implicit and explicit retention tests. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 1251–1269.
- Roskies, A. L., Fiez, J. A., Balota, D. A., Raichle, M. E., & Petersen, S. E. (2001). Task-dependent modulation of regions in the left inferior frontal cortex during semantic processing. *Journal of Cognitive Neuroscience*, 13, 829–843.
- Rundus, D. (1973). Negative effects of using list items as retrieval cues. *Journal of Verbal Learning and Verbal Behavior*, 12, 43–50.
- Sakagami, M., & Niki, H. (1994). Spatial selectivity of go/no-go neurons in the monkey prefrontal cortex. *Experimental Brain Research*, 100, 165–169.
- Samuel, A. (1996). Does lexical information influence the perceptual restoration of phonemes. *Journal of Experimental Psychology: General*, 125, 28–51.
- 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.
- Schooler, J. W., & Engstler-Schooler, T. Y. (1990). Verbal overshadowing of visual memories: Some things are better left unsaid. *Cognitive Psychology*, 22, 36–71.
- Schooler, J. W., Fiore, S. M., & Brandimonte, M. A. (1997). At a loss from words: Verbal overshadowing of perceptual memories. In D. Medin (Ed.), *The psychology of learning and motivation* (Vol. 37, pp. 293–334). San Diego: Academic Press.
- Shaw, J. S., Bjork, R. A., & Handal, A. (1995). Retrieval-induced forgetting in an eyewitness-memory paradigm. *Psychonomic Bulletin & Review*, 2, 249–253.
- Shivde, G., & Anderson, M. C. (2001). The role of inhibition in meaning selection: Insights from retrieval-induced forgetting. In D. S. Gorfein (Ed.), *On the consequences of meaning selection: Perspectives on resolving lexical ambiguity* (pp. 175–190). Washington, DC: American Psychological Association.
- Shimamura, A. P. (1995). Memory and frontal lobe function. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 803–813). MIT Press.
- Simpson, G. B., & Kang, H. (1994). Inhibitory processes in the recognition of homograph meanings. In D. Dagenbach, & T. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 359–381). San Diego, CA: Academic Press.
- Smith, A. D. (1971). Output interference and organized recall from long-term memory. *Journal of Verbal Learning and Verbal Behavior*, 10, 400–408.
- Smith, A. D. (1973). Input order and output interference in organized recall. *Journal of Experimental Psychology*, 100, 147–150.
- Smith, A. D., D'Agostino, P. R., & Reid, L. S. (1970). Output interference in long-term memory. *Canadian Journal of Psychology*, 24, 85–89.
- Smith, E. E., Adams, N. E., & Schorr, D. (1978). Fact retrieval and the paradox of interference. *Cognitive Psychology*, 10, 438–464.
- Smith, R. E., & Hunt, R. R. (2000). The influence of distinctive processing on retrieval-induced forgetting. *Memory & Cognition*, 28, 503–508.
- Smith, S. M., & Tindell, D. R. (1997). Memory blocks in word fragment completion caused by involuntary retrieval of orthographically related primes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 355–370.
- Thompson-Schill, S. (1997). Role of the left inferior prefrontal cortex in the retrieval of semantic knowledge: A reevaluation. *Proceedings of the National Academy of Sciences USA*, 94, 14792–14797.
- Tsukimoto, T., & Kawaguchi, J. (2001). *Retrieval-induced forgetting: Is the baseline in the retrieval-practice paradigm true?* Poster session presented at the International Conference on Memory, Valencia, Spain.
- Tulving, E., & Arbuckle, T. Y. (1963). Sources of intratrial interferences in paired-associate learning. *Journal of Verbal Learning and Verbal Behavior*, 1, 321–334.
- Tulving, E., & Arbuckle, T. Y. (1966). Input and output interference in short-term associative memory. *Journal of Experimental Psychology*, 72, 145–150.
- Wheeler, M. A. (1995). Improvement in recall over time without repeated testing: Spontaneous recovery revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 173–184.
- Williams, C. C., & Zacks, R. T. (2001). Is retrieval-induced forgetting an inhibitory process. *American Journal of Psychology*, 114, 329–354.