

The Role of Inhibitory Control in Forgetting Unwanted Memories: A Consideration of Three Methods

Michael C. Anderson

University of Oregon, USA

Summary. When confronted with reminders to things that we would prefer not to think about, we often attempt to put the unwanted memories out of awareness. Here, I argue that the ability to control memory is a special case of a broad class of situations thought to require executive control: response override. In such situations, one must stop a strong habitual response to a stimulus due to situational demands, a function thought to be accomplished by inhibitory processes that suppress the response, enabling more flexible, context-sensitive control over behavior. Recent behavioral studies show that inhibitory mechanisms that control overt behavior are also targeted at declarative memories to control retrieval. Recent neuroimaging findings (Anderson et al., 2004) further establish that controlling awareness of unwanted memories is associated with increased dorsolateral prefrontal cortex activation, reduced hippocampal activation, and impaired retention of the unwanted trace and that the magnitude of activation in prefrontal cortex predicts memory suppression. These findings indicate that cognitive and neural systems that support our ability to override prepotent responses can be recruited to override declarative memory retrieval, and that this cognitive act leads to memory failure. The relation between these findings and those obtained with the directed forgetting procedure is also discussed.

Key words. Inhibition, executive control, forgetting, prefrontal cortex.

Introduction

In the film *Eternal Sunshine of the Spotless Mind*, Joel suffers from the pain of a broken relationship with his former love, Clementine. To rid himself of his pain, he visits a memory clinic and arranges to have all memories of Clementine erased from his brain. The deletion procedure in-

volves presenting objects that remind Joel of Clementine while he has his brain scanned to construct a “map” of all the memories of Clementine that are in his brain. Once the Clementine map is constructed, technicians perform a memory “deletion” procedure. They go through this procedure in the convenience of Joel’s home (while he is asleep in bed), erasing each memory, one at a time, through highly focused brain damage.

Although the technology imagined in this film may be far-fetched, the film highlights an important theme that often goes unappreciated: Sometimes it is not desirable to have good memory for an experience. Like Joel, we often encounter reminders to things that we would rather not think about. Whether we are reminded of a past relationship, a lost loved one, a violent attack, a task we would prefer not to do, or something as mundane as an old telephone number when a more recent one is desired, we are frequently disrupted by a memory system that is at times too efficient in delivering memories that we do not want or need. In these circumstances, we often exert effort to put these memories out of mind, and we may wish that a memory deletion device existed. Even concentrating on a single idea or train of thought relies upon the capacity to prevent ourselves from being diverted in the different directions our mind might wander based on spreading activation. Remaining focused requires controlling the retrieval of distracting ideas. Given its broad importance, a central goal of cognitive neuroscience ought to be to elucidate the mechanisms by which such control is achieved.

This chapter reviews the work done in my laboratory examining the mechanisms underlying the control of unwanted memories. Although the findings we have obtained are far from “memory deletion,” they do indicate that people have some capacity to suppress unwanted memories. A core claim that my colleagues and I make is that this ability to control retrieval is supported by executive control processes of the sort that are widely studied in cognitive psychology and cognitive neuroscience, but targeted at declarative memories. Specifically, we argue that inhibitory control processes can be recruited to stop or override memory retrieval and thereby to exclude unwanted memories from consciousness, and that this cognitive act contributes to later memory failures for the excluded trace. By this perspective, an intimate connection exists between our many memory failures and the ability to control the direction of cognition.

The work discussed in this chapter is related to another line of research that examines the ability to control unwanted memories: directed forgetting (see Sheard & MacLeod, this volume). Although these lines of research have progressed separately, it is important to consider the relations between them, and how this may add to our understanding of the different ways in which memory control may operate. Thus, a second aim of this

chapter is to discuss alternative views of the relation between the mechanisms involved in these different experimental approaches to the inhibition of unwanted memories. In so doing, a framework is offered with the goal of fostering theoretically targeted comparisons of these approaches. By isolating the mechanisms that contribute to motivated forgetting of past experiences, we may better understand the limits of our naturally occurring efforts to achieve a “spotless mind.”

Executive Control and the Mechanisms of Retrieval

The current perspective begins with the observation that actions, once started, can usually be stopped. This fact was impressed upon me one evening while opening the kitchen window. As the window slid along its track, it pushed a cactus off the sill. My hand darted to catch the cactus. Mere centimeters from it, I stopped my hand from clutching the cactus’s needle-dense body. This timely save was made possible by my ability to stop physical action—an ability so pervasive that it goes unnoticed in daily life.

The preceding case is a classic example of a situation in which we need to stop a strong habitual response, a situation widely regarded as requiring executive control. This is sometimes referred to as response override, and is illustrated in Figure 1. In response override, one must stop a prepotent response to a stimulus, either because the response needs to be withheld or because a less common response is more appropriate. The capacity to stop or redirect action in this way is crucial. Without it, we would lose flexibility to adapt behavior according to changes in goals or to changes in the environment. We would be slaves to habit or reflex.

But how do we keep from being automatically controlled by the habitual action? One widely discussed answer is that inhibition is used to suppress the habitual response. By this view, the appearance of a stimulus activates a representation of that stimulus in memory. Activation then spreads 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 there are multiple responses, the one that achieves threshold most quickly will generally be emitted. However, if a weaker response is more appropriate, inhibition can be recruited to suppress the stronger ones. Inhibition is thought to reduce the activation level for a given response, preventing it from achieving threshold. In so doing, weaker, but more appropriate responses can be expressed, enabling flexible, context-sensitive behavior. This is known as inhibitory control.

Typical Response Override Situation

(e.g., Stroop or Go/No-Go tasks)




Fig. 1. A typical response override situation. Circles correspond to representations in memory, with lines representing associations between these representations. The stimulus is linked to two responses, one of which is stronger (prepotent), and the other of which is weaker (dotted line). Response override must occur when the organism needs either (a) to emit the weaker, but more contextually appropriate response, despite the stronger association to the prepotent response, or (b) to 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.

Given the importance of inhibitory control in managing overt behavior, one might ask whether internal actions might also be influenced by such mechanisms. 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 activate 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 a memory associated to the cue; other times, we may wish to avoid retrieval altogether either because the memory is unpleasant or because we wish to maintain focus on the cue concept. Although we often retrieve things that we do not intend, we can control this tendency; we can recollect the event we are seeking despite interference from stronger competitors, and we can stop ourselves from thinking about unwanted memories. Given these parallels

between motor behavior and retrieval, response override mechanisms may be recruited to control unwanted memories. If so, we should find evidence for inhibitory control in memory situations likely to involve response override.

By our view, strong evidence for these parallels exists. In support of this, I briefly review evidence for inhibition in memory selection and in memory stopping. Memory selection is required when our goal is to recall an event or fact from 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 and wish to prevent an associated memory from entering awareness. In both cases, attempts to limit the influence of distracting memories have been found to impair later retention, highlighting an important link between forgetting and the control of retrieval.

Selective Retrieval and Retrieval-Induced Forgetting

The need to select a weaker response to a stimulus in the face of interference from a prepotent competitor finds a parallel in memory in the situation of selective retrieval. Here, the aim is to recall a target memory when given one or more cues. Typically, a cue will be associated with other memories as well—and some of those other memories may be more associated to the cue than is the target item. It is well known that when multiple traces are associated to the same cue, they compete for access to consciousness (see Anderson & Neely, 1996; Postman, 1971, for reviews). This form of competition presents a problem of control because the cue cannot be relied upon to access the target—in fact, the presence of a strong competitor could perpetually divert us from that target. If inhibitory control is recruited to override prepotent responses, then inhibition might also be used to override prepotent memories. To the extent that inhibition persists, situations demanding the selective retrieval of a target should induce lasting memory impairment on competitors. Thus, the act of remembering should cause forgetting of related memories.

Over the last decade, my colleagues and I have explored this prediction with a procedure that we refer to as the retrieval practice paradigm (Anderson, Bjork, & Bjork, 1994). In the typical study, subjects encode lists of category-exemplar pairs (e.g., fruit – orange, drink – 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 several times to increase the effect of retrieval practice on related items. After a delay, subjects are tested on




Fig. 2. 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 (6 items are usually studied in 8 categories). 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.

all studied exemplars. Performance on this category cued recall 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). Figure 2 illustrates typical findings. Not surprisingly, recall of the practiced exemplars was improved relative to performance on baseline items. More importantly, recall for the unpracticed exemplars from the practiced categories (e.g., banana) was worse than that for the items from baseline categories (e.g., drinks). Thus, remembering some items during retrieval practice led subjects to forget related items on a delayed test. We refer to this finding as *retrieval-induced forgetting* (Anderson et al., 1994) to highlight the central role that retrieval plays in generating the effect. Retrieval-induced forgetting is consistent

with the view that inhibitory control is recruited to combat interference during retrieval, with inhibition manifesting as recall impairment for competitors on the final test.

Although the basic finding of retrieval-induced forgetting is compatible with inhibition, other mechanisms can explain this effect as well, including McGeoch's (1942) classical response competition theory of interference. According to this theory, the likelihood of recalling a target should decrease either when a new response gets associated to the cue used to retrieve it, or when an existing alternative response is strengthened. In either case, the target will suffer increased competition from the alternative response. These competitive dynamics have become formalized in several memory theories that posit relative strength rules of retrieval (e.g., Raaijmakers & Shiffrin, 1981; J.R. Anderson, 1983). In these models, the probability of recalling a target is determined by that item's association to a cue, relative to the strengths of association of all items related to that cue. When an alternative response is strengthened, say by retrieval practice, the relative strength of all nonpracticed items declines. Later, when the subject tries to recall the target, the strengthened competitor will have a retrieval advantage, leading it to intrude so persistently that the subject abandons efforts to recall the unpracticed exemplars (see also, Rundus, 1973).

This approach does not require inhibition; rather, practiced items become so strongly linked to the practice cue that they block other exemplars. This account is plausible, given the strengthening that practiced items enjoy (but see the later section in this chapter on properties of retrieval-induced forgetting). Other mechanisms may also contribute to retrieval-induced forgetting. For example, retrieval practice may damage the association linking the category to the affected exemplar or instead may alter 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 impairment, see Anderson & Bjork, 1994). Although it might seem difficult to distinguish these alternatives, focused empirical research has yielded evidence for properties of retrieval-induced forgetting that favor the inhibition view.

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 (see Anderson, 2003, for a review). First, several findings demonstrate that strengthening practiced

items does little to impair the recall of related competitors provided that all sources of retrieval-induced forgetting are eliminated from the strengthening process and from the measurement of impairment. For instance, retrieval-induced forgetting appears to be *recall-specific*: Retrieval practice impairs the delayed recall of competing items, but the same number of repeated study exposures does not (provided that output interference is controlled on the final test), even though the two practice procedures strengthen practiced items to the same degree. The fact that strengthening can occur with little associated impairment suggests that retrieval practice, not strengthening, is responsible for the effect. Consistent with this, when retrieval practice is performed, the amount of impairment often has no relation to the amount of strengthening observed on practiced items—that is, retrieval-induced forgetting appears to be *strength-independent*.

Impairment does appear to be *interference-dependent*, however: Whether retrieval practice impairs a related item depends on whether the item causes interference during retrieval practice. So, for instance, high frequency exemplars of categories suffer retrieval-induced forgetting whereas low frequency exemplars do not; similarly, the dominant meanings of asymmetric homographs suffer significant retrieval induced forgetting, whereas the subordinate meanings do not. Interference-dependent impairment is exactly what one would expect if inhibitory control is recruited to override retrieval of distracting competitors.

Finally, retrieval-induced forgetting exhibits a crucial theoretical property that is difficult for traditional associative interference accounts to explain: *cue-independence*. Cue independence refers to the tendency for retrieval-induced forgetting to generalize to novel cues other than those used to perform retrieval practice. Retrieval practice on Fruit-Orange not only impairs the later recall of Banana when it is tested with the cue Fruit, but also when it is tested with a novel, independent cue such as Monkey B____. These findings argue against interpretations of retrieval-induced forgetting such as associative blocking or cue change, which posit mechanisms that are specific to the cues guiding retrieval practice. These findings are, however, consistent with the idea that competing memories themselves are suppressed by an inhibitory process in order to retrieve the retrieval practice target.

Taken together, the foregoing properties indicate that the impairment underlying retrieval-induced forgetting is unlikely to be produced by traditional associative interference mechanisms. Rather, it is likely to reflect the action of an inhibitory control process acting to override unwanted retrievals of competitors in memory, helping to achieve selective memory retrieval. This supports the view that selective memory retrieval may be re-

garded as a special case of response override arising in long-term memory retrieval.

Stopping Memory Retrieval

Response override is also involved when we need to stop a response. In retrieval, the ability to stop could prove useful in preventing a memory from entering consciousness. 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 a movie may lead us to stop the natural progression from cues to memories. Other times, we may wish to focus on a thought without letting our mind wander. Can inhibitory mechanisms be engaged to serve these goals?

Anderson and Green (2001) looked at this issue by examining how stopping retrieval affected the memories that were to be retrieved. To study this, they developed a 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 humans (e.g., Casey et al., 1997; Garavan et al., 1999) and monkeys (e.g., Sakagami & Niki, 1994). In one version of the this task, letters are presented one at a time and subjects 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 require a button press so that, when an X occurs, subjects have difficulty withholding their 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; hereafter 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 control over retrieval. For most of the trials, the task was the same as it had been during training—to recall and say aloud the associated word as quickly as possible at the sight of its cue. For certain cues, however, subjects were admonished to avoid thinking of the response. It was emphasized that it was not enough to avoid *saying* the response word—it was crucial to prevent the memory from entering *conscious awareness* at all. Thus, subjects had to override not only a vocal 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 controlled consciousness. However, if inhibitory control was 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, and asked to recall the response for each. As predicted, forgetting occurred: Response words that subjects excluded from 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 became (see Figure 3a). Interestingly, avoided words were harder to recall even though subjects had seen as many as *16 reminders* (i.e., cues) during the think/no-think phase. Normally, reminders facilitate memory, much as they did for the items to which subjects continued to respond (Figure 3a).

Anderson and Green (2001) further established that this impairment was cue independent, echoing the results of Anderson and Spellman (1995): Forgetting occurred regardless of whether subjects were tested with the studied cue word (e.g., ordeal) or with a novel cue never studied in the experiment (e.g., insect r_____ for roach; Figure 3b). This argues that the forgetting is not caused solely by associative interference; rather, impairment reflects suppression of the excluded memory itself. In a control 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, indicating that the recall deficits 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 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 to regulate awareness of intrusive memories. In particular, whenever the environment 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 (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,




Fig. 3. Final recall performance in Experiment 1 of Anderson and Green (2001) using the think/no-think procedure. The 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 (suppress), or tried to recall it (respond). The top panel represents final recall performance when tested with the originally trained retrieval cue (i.e., the “Trained probe”), whereas the bottom panel represents final recall performance when tested with a novel, independent, extralist category cue.

2001; Bjork et al., 1998; Conway et al., 2000; Deprince & Freyd, 2001; Myers, Brewin, & Power, 1998; see Golding & MacLeod, 1998, for a review of directed forgetting). More broadly, these findings may be related to cognitive situations in which people must “deselect” unwanted aspects of the representation of a stimulus to attend to task relevant attributes (see, e.g., Rajaram, Srinivas, & Travers, 2001; Rajaram, this volume; see also, Humphreys, this volume).

Neural Systems Underlying Voluntary Memory Suppression

The foregoing work on selective retrieval and retrieval stopping indicates functional parallels between controlling retrieval and overriding prepotent responses. This suggests the intriguing possibility that the ability to control unwanted memories may in part rest on neural systems essential for controlling overt behavior. More direct evidence for this relation might be observed if more were known about the anatomical systems that support memory control. Recently, we have used neuroimaging to identify the neural systems underlying this ability. If memory control and response override are related, one might expect stopping memory retrieval to recruit neural systems known to be involved in overriding prepotent responses to control structures involved in memory.

Research on the neural basis of executive control and declarative memory indicates that at least two brain regions may play important roles in the neurobiological basis of memory control: the hippocampus and the lateral prefrontal cortex. The hippocampus is essential for declarative memory formation (Squire, 1992), and increased hippocampal activation is associated with the subjective experience of consciously recollecting a recent event (e.g., Eldridge et al., 2002). Memory suppression requires people to stop retrieval to prevent conscious recollection. Lateral prefrontal cortex is involved in overriding prepotent motor responses (e.g., Garavan et al., 2002; Menon et al., 2001; Aron et al., 2003), switching task set, and combating interference in a range of cognitive tasks (e.g., Aron et al., 2004; Knight et al., 1999; Shimamura, 2000). We hypothesized, therefore, that people suppress consciousness of unwanted memories by recruiting lateral prefrontal cortex to disengage hippocampal processing that supports recollection.

Anderson et al. (2004) recently examined this fronto-hippocampal hypothesis using the think/no-think procedure. Subjects first learned pairs of words. They then received trials in which they were given the first member of a word pair and asked (if the word appeared in green) to think of its corresponding response (Respond condition) or (if the word appeared in red)

to suppress awareness of the response (Suppression condition). Subjects performed this task while being scanned in an event-related functional magnetic imaging design. After this phase was over, scanning ended, and memory was tested for all of the word pairs subjects had studied. Replicating prior work, subjects recalled significantly fewer suppression than baseline items, showing that suppression had occurred.

To examine the neural systems underlying suppression, we contrasted activation during Suppression and Respond trials of the think/no-think phase. In both trial types, subjects had been presented with a cue word for 4 seconds that had previously been associated with a response word; the only difference was that the cues presented in red had directed subjects to suppress consciousness of the unwanted memory, whereas those presented in green had directed subjects to recall and maintain the corresponding response. Several remarkable findings were obtained. First, a network of regions was more active during suppression than during active retrieval, including bilateral dorsolateral and ventrolateral prefrontal cortex, and anterior cingulate cortex. This network overlaps strongly with networks involved in motor response suppression tasks (such as go/no-go), even though no motor responses were required. These findings support the possibility that neural systems involved in overriding prepotent responses are recruited to control retrieval. More generally, they indicate that regulating consciousness of unwanted memories is an active process, and is not accomplished by simply failing to engage retrieval mechanisms.

Interestingly, Anderson et al. (2004) also found that suppressing awareness of a memory significantly reduced activation in the hippocampus bilaterally, relative to retrieval. Because activation in the hippocampus has been linked to conscious recollection, this finding suggests that subjects can control awareness of past experiences by strategically disengaging activation in the hippocampal memory system that might otherwise support conscious recollection.

One might wonder to what extent the suppression regions observed in the overall analysis are functionally involved in suppressing unwanted memories. More compelling support for this role could be obtained if it could be shown that these regions predicted later memory suppression effects. We evaluated this by using regression to examine which brain regions predicted individual differences in memory inhibition. Crucially, activation in dorsolateral prefrontal cortex and lateral premotor cortex, regions often observed in go/no-go task performance, predicted subsequent memory impairment for suppression items (Figure 4a). The more activation there was in these regions, the more suppression subjects showed (Figure 4b). We further showed that activation in the hippocampus predicted later memory failures for suppressed items and that these variations

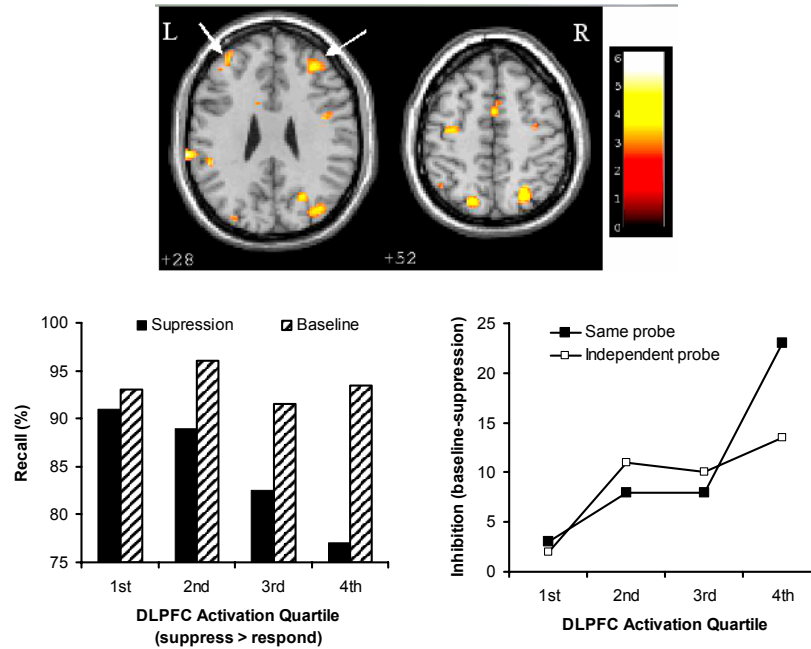


Fig. 4. Relation of suppression-related activations to memory inhibition. (a) Regions for which activation during suppression trials predicted differences in below-baseline inhibition ($n = 24$). White arrows highlight DLPFC regions from the regression analysis that also predict hippocampal activity for suppression items. (b) Memory inhibition effects for four subject groups matched for counterbalancing manipulations, differing in degree of activation in right DLPFC. Note that increasing activation in DLPFC predicts reduced suppression performance, but leaves baseline performance unaffected. (c) Memory inhibition effects (Baseline – Suppression) in four DLPFC groups, separated by test type.

were correlated with activation in the dorsolateral prefrontal cortex. These findings suggest that dorsolateral prefrontal cortex interacted with medial-temporal lobe structures to attenuate activity, reducing recollection and disrupting retention.

This work provides a strong indication that some of the neural systems involved in overriding prepotent responses may also be recruited to terminate internal actions such as retrieval. Such systems appear to be targeted at medial-temporal regions that support declarative memory rather than motor representations. Because ideas and memories are brought into consciousness by retrieval, the capacity to stop retrieval provides a specific

cognitive and neurobiological foundation for how human beings regulate consciousness of unwanted memories. This work further supplies a model for how motivated forgetting occurs that may be applicable to understanding the adaptation of memory in the face of traumatic experiences. By integrating this model with insights derived from other empirical approaches, we can more fully characterize the variety of ways in which people may regulate awareness of unwanted memories.

Forgetting Unwanted Memories: A Comparison of Methods

So far in this chapter, I have focused on research using the retrieval practice and think/no-think paradigms. Other paradigms have been used to study the inhibition of unwanted memories, however. In this final section, I will discuss the relation of the current procedures to an important method used to study intentional forgetting: Directed forgetting. By considering the relation between these paradigms, I hope to encourage direct comparison of the methods, and foster principled theoretical development. I begin with a brief review of different methods used to study directed forgetting, along with some preliminary considerations of the nature of inhibitory effects in directed forgetting studies. I then offer a simple framework within which the paradigms for studying inhibition effects may be considered.

Directed Forgetting: the Phenomenon and its Interpretations

The term *directed forgetting* refers to impaired memory arising from an instruction to forget the unwanted material. Such impaired memory has been readily observed in two experimental procedures, known as *item method* and *list method* directed forgetting (see Sheard & MacLeod, this volume; see also Golding & MacLeod, 1998). In the item method, subjects typically view a series of words, each followed by an instruction cuing them to remember the item for a later test, or to forget it. After all of the words have been presented, memory is tested with either a recall or a recognition test. The item-by-item instruction manipulation yields a sizeable recall deficit for “forget” items relative to remember items that occurs on both recall and recognition tests. Directed forgetting in studies using the item method is thought to reflect selective rehearsal of the to-be-studied words. For instance, subjects may rehearse words shallowly (e.g., by subvocal repetition) until they receive the cue, at which point they either (a) stop rehearsing the word in the case of the forget instruction, or (b) elaboratively

encode the word in the case of the remember condition. This view is consistent with the impaired performance of forget items on recognition tests. If correct, this interpretation suggests that item method directed forgetting is more properly regarded as an encoding effect, rather than as evidence for inhibition.

The list method differs from the item method mainly according to when the instruction to forget is presented. In the list method, subjects receive an instruction to forget or to remember the studied items only after a long list (typically 10-20 words) has been encoded. The instruction, furthermore, is a surprise and so subjects are likely (prior to the instruction) to extend their best efforts to encode the words. This feature of the procedure is thought significant because it makes it less likely that differential encoding of first-list items could underlie recall deficits arising from the forget instruction (however, see Sheard & MacLeod, current volume). After the forget or remember instruction is given, subjects study a second list. Once both lists have been presented, memory is tested. The test may require recall of the first list, the second list, or both lists.

The list method often yields three effects that typify this form of directed forgetting: (a) impaired recall for the first list of items when subjects are instructed to forget the first list, compared to when they are to remember it (i.e., directed forgetting *costs*); (b) improved recall for the second list of words when subjects are instructed to forget the first list, relative to when they are to remember the first list (i.e., directed forgetting *benefits*); and (c) superior memory for second-list words compared to first-list words in the forget-instructed group. These effects are generally restricted to recall tests, with little effect observed in recognition.

This pattern has led investigators to attribute list method directed forgetting to retrieval inhibition. By this view, items on the first list are inhibited by the instruction to forget, but remain available in memory, as evidenced by intact performance on recognition tests. If this analysis is correct, then list method directed forgetting is more likely to have mechanisms in common with the retrieval-practice and think/no-think procedures than would item method directed forgetting. However, even with the list method, several restrictions should be placed on what constitutes evidence for directed forgetting. Two of these are discussed next.

Comparing Recall Performance Across Lists 1 and 2 is Not a Good Measure of Directed Forgetting

Many directed forgetting studies simplify the design by eliminating the control group in which subjects are instructed to remember the first list before proceeding to the second list. These studies have one primary condi-

tion—the forget group—and recall performance is compared for items on the first (the forget list) and second (the remember) lists. The difference in recall across these lists is taken as a metric of directed forgetting because the two lists are believed to differ only by the instruction given after the list is encoded. Superior recall of second list items is thought to reflect a mixture of the costs on first list items and the benefits on second list items.

The method of comparing list-1 and list-2 recall is fraught with difficulties, and should not be taken as retrieval inhibition. The assumption that the two lists differ only by the nature of the instruction is not correct, leading many other factors to get mixed into the estimate of inhibition. Several such differences exist. First, the second list is studied more recently, and this difference surely confers a recall advantage on the second list that has little to do with directed forgetting. Second, by the time the second list is presented, subjects have been exposed to the task, and to the nature of the stimuli that they are likely to receive. This difference may alter subjects' approach to the second list, as suggested by classic research in verbal learning on warm-up increment and learning-to-learn (Postman, 1971). Thus, recall may improve on the second list because subjects simply get into the rhythm of the task, or, alternatively, develop more well-tuned strategies for encoding items—strategies that may be qualitatively different from those used in the first list. Finally, the second list may cause retroactive interference, even in the absence of any instruction to forget the first list.

Collectively, these variations across list-1 and list-2 are confounding variables that compromise interpretation of differences in recall arising from the forget instruction, and so these differences should not be taken as evidence for directed forgetting. These confounds are avoided when directed forgetting is estimated by comparing the first list in the forget group to a control group instructed to remember those same items. I view this list-1 comparison (i.e., the cost comparison) as the only pure measure of retrieval inhibition, to the extent that such a process contributes to directed forgetting (see Sheard & MacLeod, this volume, for a discussion of alternative mechanisms that may also contribute to list-1 costs).

Studies of Directed Forgetting that Do Not Control Output Order Conflate Directed Forgetting with Output Interference

Even when a study estimates directed forgetting with a cost comparison, differences do not necessarily reflect the action of the forget cue. Differences may instead reflect differential output interference across the forget and remember conditions. Consider a directed forgetting experiment in which subjects are asked on the final test to recall both the first and second

lists in any order. If the recall of list-1 items is worse after a forget instruction than after a remember instruction, does this mean that list-1 was suppressed as a result of the forget instruction? Not necessarily. The forget subjects may have remembered fewer list-1 items because they were biased to recall list-2 items earlier.

Such a bias could arise in several ways. First, because a forget instruction on list-1 typically improves recall for list-2 items, this enhanced accessibility (relative to the remember group) may lead list-2 items to be recalled earlier in the forget group. Enhanced recall of list-2 items might arise because they are better encoded when subjects do not have to try to remember list-1 items as well. Second, the instruction to forget the first list may bias subjects to begin their recall with list-2 items, merely because they are implicitly deemed more important. Subjects in the remember group are responsible for remembering both lists and would have neither of these biases. As such, the groups may differ in the propensity to recall list-2 items early, yielding an output order bias against list-1 items that could produce recall deficits in the forget group. Because output interference is considered a form of retrieval-induced forgetting, this method of measuring directed forgetting mixes retrieval-induced forgetting and “true” directed forgetting effects.

There are several easy solutions to the foregoing problem. Subjects may be directed to recall only the list-1 items in both the forget and the remember group or, instead, to recall list-1 items followed by list-2. Directing subjects to recall first list items right away eliminates output interference from second list items and matches this factor in the remember and forget groups. Thus, any remaining differences should reflect the forget instruction and not output interference. It should be noted, however, that although estimating directed forgetting with this method provides a more theoretically focused measure of retrieval inhibition, the less controlled method provides different information. In real life settings, the total negative impact of an effort to forget may be determined by the “true” inhibitory effect of directed forgetting, compounded by retrieval biases and the accumulating output interference effects they produce. Nevertheless, when the goal is to isolate properties of the inhibition mechanism, the more focused procedure for estimating inhibition should be used.

Views on the Relation of Directed Forgetting to the Current Phenomena

Given the preceding constraints on evidence, the relation between directed forgetting, retrieval-induced forgetting, and memory impairment produced

by the think/no-think procedure can be understood in several ways. Some have argued for a distinction between controlled and automatic forms of inhibition. I will discuss this perspective, and some of the evidence thought to support it. I will then propose an alternative view that rests on two factors: the intention to forget, and the level of representation affected by inhibition. I will end by considering the implications of this framework for proposals concerning the mechanisms underlying directed forgetting.

Controlled versus Automatic Inhibition

Some have argued that retrieval-induced forgetting and directed forgetting differ in the level of executive control involved in producing inhibition (e.g., Conway & Fthenaki, 2003). Several considerations motivate this proposal. First, directed forgetting requires an intentional effort to forget, whereas retrieval-induced forgetting does not. Because goal directed cognition generally requires cognitive control, explicit instructions to forget should place demands on controlled attention. Second, some evidence suggests that directed forgetting requires attention to be performed properly. Dividing attention during the encoding of the second list disrupts directed forgetting on the first list (Conway, Harries, Noyes, Racsma'ny, & Frankish, 2000). Directed forgetting also appears to be diminished in populations thought to suffer deficits in executive control, such as in the elderly (Zacks, Radvansky, & Hasher, 1996), young children (Harnishfeger & Pope, 1996), and frontal-lobe damaged patients (Conway & Fthenaki, 2003). Together, these findings suggest that directed forgetting requires controlled inhibition. Although these arguments have not been formally extended to the think/no-think procedure, the fact that this procedure instructs subjects to exclude a memory from awareness would seem, by the same logic, to require controlled processing that ultimately leads to inhibition.

Retrieval-induced forgetting, by contrast, might seem to require less control. In this procedure, no instruction to forget is given; rather, forgetting is a by-product of retrieving related material. Because these effects are unintentional, they may be produced by automatic processes. Automatic inhibition mechanisms can certainly be envisioned for retrieval-induced forgetting: Retrieving targets may inhibit competitors by means of automatic lateral inhibitory connections, for example. Consistent with the automaticity view, several studies have found intact retrieval-induced forgetting in populations thought to be deficient in executive control. Moulin, Perfect, Conway, North, Jones, and James (2002) found that older adults with and without Alzheimer's disease showed robust retrieval-induced forgetting, contrary to what these authors expected if these populations had

deficits in controlled inhibition. Conway and Fthenaki (2003) found that frontal patients showed diminished directed forgetting; however, frontal patients exhibited significant, though reduced retrieval-induced forgetting. Because the prefrontal cortex is thought to support executive control, these findings appear compatible with the view that retrieval-induced forgetting is produced by a more automatic form of inhibition.

Although the foregoing view is plausible, there are problems with the arguments for it. The major difficulty arises from the fact that all of the preceding studies used a final memory test that provided subjects with the same cue used to perform retrieval practice. For instance, when subjects performed retrieval practice on several fruits, later recall was assessed with “Fruits” as the cue rather than a separate cue designed to test the accessibility of the inhibited items. Because of this, the measure of retrieval-induced forgetting mixes impairment arising from suppression with interference from the practiced items. Thus, when subjects recall Fruits, they may perseveratively recall the stronger practiced items, blocking access to competitors.

These two sources of impairment—blocking and suppression—should be differentially affected by attentional deficits (and also by dividing attention). For example, deficient attentional inhibition should reduce the contribution of suppression to retrieval-induced forgetting, decreasing the effect. In contrast, that same deficiency should increase vulnerability to interference from the practiced items, exaggerating the effect. Thus, even if frontal lobe or Alzheimer’s patients had no capacity for attentional inhibition, they should show sizeable retrieval-induced forgetting to the extent that these deficits render subjects unable to combat interference from practiced competitors on the final test. Indeed, this vulnerability to interference (the tendency for frontal patients to perseverate strong responses) forms the basis for the inhibitory deficit hypothesis of frontal lobe function. Thus, no conclusions can be inferred about inhibitory deficits when the independent probe method is not used. There is little reason, at present, to conclude that retrieval-induced forgetting does not require cognitive control.

Independent of these empirical arguments, one may question whether the lack of intention to forget competing items in the retrieval-induced forgetting procedure should be equated with a lack of cognitive control. Although this seems plausible at first, I will argue in the next section that such a linkage is entirely unnecessary and perhaps incorrect.

The Flexible Control Hypothesis: An Alternative Framework

Retrieval-induced forgetting, impairment in the think/no-think paradigm, and directed forgetting may not stem from different types of inhibition varying in automaticity. Rather, they may reflect a common inhibitory process that is flexible in two important ways. First, it seems plausible that the inhibition mechanism may be *goal general*—that is, it can be recruited for different cognitive goals, including memory retrieval, working memory maintenance, selective attention, avoidance of unwanted thoughts, and intentional forgetting. The idea that inhibition is goal general is implicit in many proposals concerning inhibition as a controlled process. Nevertheless, it is worth emphasizing here, because it provides an important proposition necessary to our account of the relations between these phenomena. Second, the inhibition mechanism may be *representation general*—that is, once recruited, it can be targeted at different types of representation. Inhibition can be targeted at episodic or semantic memories, and at memories varying in content. Importantly, in the current proposal, inhibition can be targeted at different levels of representation. Inhibition may be targeted at individual items within a context, or at a global contextual representation. Here again, this flexibility is implicit in the notion of controlled inhibition. The idea that inhibition can be targeted at individual items or contexts will be featured in our account of the differences between inhibitory phenomena. These considerations suggest that it may be worthwhile to evaluate the ways in which memory inhibition tasks differ in their goals and in the targets of inhibition. We consider each of these dimensions in turn.

Differences in Goals, Not in Control

All of the tasks discussed in this chapter differ in the goals for which inhibition is recruited. In retrieval-induced forgetting, the goal is to recall target items, given (in many cases) a category plus a letter stem cue. No instruction is given to forget competing exemplars, nor are competitors mentioned. Whatever inhibition occurs thus does not arise from an explicit goal to forget competitors, but rather from a goal to retrieve target items. In directed forgetting, however, the putative inhibition process is initiated by the “forget” instruction itself, although effective inhibition is also believed to rely on encoding new information in a second list. Whatever inhibition occurs on to-be-forgotten items is tied to an explicit goal to forget that makes mention not of individual items per se, but of the entire first list. Finally, in the think/no-think procedure, inhibition takes place when subjects exclude the unwanted memory from consciousness. No instruction to forget is given, though the to-be-inhibited item is mentioned be-

cause the task is to exclude that item from awareness. Whatever inhibition accrues arises from a goal that concerns the target, though the goal is not to forget.

The preceding analysis illustrates how inhibition may be recruited in support of different goals. One might wonder, however, whether the inhibition involved in each case is the same. According to the controlled/automatic distinction, for example, the inhibition underlying directed forgetting and effects in the think/no-think paradigm differs qualitatively from the form of inhibition at work in retrieval-induced forgetting. I believe, however, that this proposal confuses two similar, though separable dimensions: the explicitness of the goal to forget, and the presence of executive control. According to the flexible inhibition hypothesis, controlled inhibition may be recruited in service of our goals, regardless of whether these goals make reference to forgetting, as long as there is a control problem. For instance, the retrieval-practice paradigm makes no reference to forgetting, but isolating the target during retrieval practice may require cognitive control to push distracting competitors out of mind. Similarly, the goal of keeping an unwanted memory out of awareness makes no reference to forgetting; nevertheless, the cue to which subjects attend activates the memory subjects are avoiding, requiring inhibitory control. Thus, although instructing subjects to forget items may entail controlled inhibition, the lack of intention in retrieval-induced forgetting does not imply an absence or even a reduction of controlled inhibition. The working assumption is that all of the paradigms discussed here (to the extent that each involves inhibition) make use of a common controlled inhibition process.

Differences in Representational Target of Inhibition

Inhibition may also be targeted at different types and levels of representation. In directed forgetting, the instruction to forget does not make reference to individual items, but rather to the entire first list. Although such an instruction might be achieved in many ways, inhibition may target the list-1 context rather than individual items. By this view, each item on the first list is encoded with the list-1 context. The context could be regarded as a discrete “list” concept or a collection of features that permits discrimination and source recollection. If the contextual representation were to be suppressed, it could reduce the accessibility of all items in this list, even though no individual item was suppressed. If subjects construct a new context representation for the second list, proactive interference from list-1 should be minimized during the encoding of list-2. Furthermore, list-1 recall should suffer if subjects use a contextual representation at test that is

favorable to list-2 items. Thus, the immediate cause of list-1 impairment may be the use of inappropriate contextual cues, although this ultimately may arise from suppression of the list-1 context.

The retrieval-practice paradigm presents a situation more likely to favor the suppression of individual items rather than of the global list context. Because the retrieval practice targets are to be retrieved from the original study list context, suppressing that context would be an ineffective solution to combating interference from related items on that list. Even if the first list context were suppressed, presenting a category name as a retrieval practice cue would prompt interference from other studied exemplars, based purely on priming. Thus, to reduce interference, inhibition must be recruited on a trial-by-trial basis, in response to intrusions of related items. If so, inhibition should affect individual items, rather than the global context (though some effects of contextual shift between the list-1 and retrieval practice phases should also be considered). Similar arguments can be made about the inhibition that arises in the think/no-think paradigm. Thus, inhibition in the directed forgetting paradigm may influence a different level of representation than is affected in the retrieval practice and think/no-think paradigms, even if a common inhibition mechanism is involved.

The functional properties of retrieval-induced forgetting and of directed forgetting support the distinction between context-level and item-level inhibition. Research on retrieval-induced forgetting suggests an item-level effect. That competitors are recalled worse than baseline items from the same study list indicates that the effect cannot be a general suppression of the study-list context. Moreover, the generalization of impairment to novel cues and the observation of impairment on cued recall, recognition (e.g., Hicks & Starns, 2004), and lexical decision tests (Veling & van Knippenberg, 2004) all support the claim that accessibility of individual items has been diminished.

Research on directed forgetting, by contrast, indicates that forget instructions impair recollection, without disrupting the items themselves. For example, the same directed forgetting procedure that produces forgetting in free recall yields priming of inhibited items on indirect tests (e.g., Basden, Basden, & Gargano, 1993). When subjects perform a fame-judgment task for a list of names, names that appeared on a to-be-forgotten list are more likely to be judged famous than names in a to-be-remembered list, even when subjects are admonished not to judge studied items as famous (Bjork & Bjork, 2003). These findings suggest that items impaired by directed forgetting remain primed. These items are impaired, however, when they have to be accessed from the temporal context, an impairment that can be alleviated when the context is provided. Subjects exhibit im-

paired source memory for list-1 items even when item recognition is intact (e.g., Geiselman, Bjork, & Fishman, 1983). However, when some items from a list are given as cues for the remaining items (Goernert & Larson, 1994), or when some items are re-exposed in a recognition test (Bjork, Bjork, & Glenberg, 1973), accessibility of list-1 item improves, as indicated by diminished costs and diminished benefits, respectively. Finally, when directed forgetting instructions are replaced by instructions to induce a new mental context at the outset of a second list, many of the characteristics of directed forgetting are created (Sahakyan & Kelley, 2002). These findings indicate that list method directed forgetting operates on context rather than on items (see also Kimball & Bjork, 2002, for relevant evidence). Although this has been interpreted as evidence against inhibition (Sahakyan & Kelley, 2002), the context shift account is consistent with inhibition at a different level of representation.

Categorizing Inhibition Tasks by Goal and Target

The preceding analysis suggests that categorizing tasks by goal and by the target may provide a useful way of viewing the relations among paradigms. Figure 5 illustrates one organization of tasks along these dimensions. Rows in this chart represent goals for which inhibition may be recruited, ordered by the degree to which subjects possess an intention to forget. Columns represent tasks that differ by the representation affected, with tasks that target items and contexts represented in the left and right columns, respectively. In this chart, retrieval-induced forgetting is unintentional and affects items, whereas directed forgetting is intentional and affects the list-level (see Kimball & Bjork, 2002, for related arguments). Forgetting in the think/no-think procedure, by contrast, has an intermediate degree of intention. Subjects are asked to exclude the unwanted memory from consciousness. The instructions make no reference to forgetting, however, and are ambiguous from the subjects' point of view. Some subjects may interpret the directions as instructions to forget, whereas others may interpret them as instructions to keep the item out of awareness momentarily. The think/no-think procedure also clearly affects individual items.

This scheme achieves several functions. First, it permits categorization of inhibition tasks, highlighting ones that may have similar properties. For example, part-set cuing inhibition and output interference are cases in which inhibition (if involved) leads to unintentional forgetting. Both tasks also affect item representations in most studies. These forms of impairment therefore might have similar properties, if their main features are captured by the dimensions under discussion. Second, the scheme encourages

Degree of Intentionality	Locus of Effect	
	Item Level	Context Level
Low	Retrieval Induced Forgetting Part Set Cuing	Context Switch Sahakyan & Kelley (2002)
Intermediate	Think/No Think Proactive Interference	Retroactive Interference
High	?	List Method Directed Forgetting

Fig. 5. A chart categorizing different tasks thought to involve inhibition, according to two dimensions: the level of representation at which inhibition acts (individual items, list context), and the extent to which subjects have an explicit intention to forget a memory. Some tasks (e.g., retroactive interference) are difficult to categorize neatly, because they may have multiple effects contributing to them, and some cells are not, as yet, represented by an existing task (e.g., the lower left cell).

reclassification of tasks that might not be thought to require inhibition. For instance, Sahakyan and Kelley's (2002) manipulation of mental context is thought to impair memory without inhibition. However, this task could be unintentional forgetting due to inhibition operating at the level of context. By this view, instructions to shift out of a context and into a new one constitute a requirement to suppress a context and to replace it with a newly retrieved one. This task does not instruct subjects to forget, but may nonetheless recruit inhibition to achieve the shift. If so, one might expect list-method directed forgetting and Sahakyan and Kelley's context shift manipulation to have similar properties. Third, this scheme highlights paradigms that do not exist, suggesting directions for novel research. For example, the bottom left cell—the intentional forgetting paradigm that operates on items—arguably does not exist. One might categorize item method directed forgetting this way, but the item method is thought to reflect differential rehearsal rather than inhibition. Could an inhibitory item level directed forgetting procedure be devised? One might adapt the think/no-think procedure to require subjects' to forget the suppression items, rather than to avoid thinking about them. Such a procedure could yield different results, which may be important in understanding the role of intention in suppressing unwanted memories.

Not all tasks will neatly fit into one of the foregoing cells, and depending on how a task is done, a given paradigm may be categorized differently. Consider the A-B, A-D retroactive interference paradigm. On the one hand, forgetting may reflect item-specific suppression of first-list responses (A-B responses). On the other hand, acquiring a second list may suppress the first-list context, inducing a list-wide reduction in performance. Both factors have been hypothesized to play a role in retroactive interference (see Postman, 1971). Retroactive interference is also another example in which the intention to forget is intermediate or a least variable. Subjects may or may not believe that they can forget the first list, and such variations may make RI more or less like directed forgetting.

Part-set cuing can also be characterized in different ways. Some studies using categorized word lists have manipulated the number of category cues provided to subjects on the recall test (rather than providing all of the category names and varying the number of exemplar cues). Robust part-set cuing is found for the remaining categories (see Nickerson, 1984, for a review). However, this type of part-set cuing affects the accessibility of uncued category names themselves, rather than of individual items. Thus, inhibition is acting at a different level of organization (the category list level) than is typically affected by exemplar cuing. Under these circumstances, it is unclear whether part-set cuing should be considered an item- or a list-level phenomenon, as access to whole categories is being affected. That the current framework prompts consideration of these questions suggests it may be useful in encouraging an integrated understanding of memory inhibition paradigms.

Summary

In the beginning of the chapter, selective retrieval and retrieval stopping were presented as two situations that require executive control to suppress unwanted memories. Other perspectives are possible, however. By one view, retrieval-induced forgetting (and other tasks like part-set cuing) differ from directed forgetting and the think/no-think paradigm in their reliance on automatic versus controlled forms of inhibition respectively. Although this is plausible, an alternative view has been proposed here: the flexible control framework. By this view, inhibition can be recruited for many goals, and targeted at different types of representation. Given these types of flexibility, controlled inhibition may be involved in both intentional and incidental forgetting tasks as long as the demand for control is present in each. I have offered a classification of these paradigms, and others, in terms of whether inhibition is intentional, and whether it is tar-

geted at individual items or temporal context. Although it is unclear whether intentionality matters to the characteristics of inhibition, this dimension differentiates the paradigms, and is separable from whether executive control is involved. The available evidence suggests, however, that the level of representation affected is an important dimension that dictates the properties of forgetting. By attending to these dimensions, we may be able to achieve a better understanding of the level of control involved in different inhibition tasks, and of the relation of these phenomena to one another.

Concluding Remarks

The need to control memory pervades daily life. The need is captured vividly by a moment with which we are all familiar. This moment occurs when, by chance, we encounter a cue in our environment that reminds us of an unpleasant past event—a brief flash of experience and feeling that is rapidly followed by an attempt to exclude the unwanted trace from awareness. We, in essence, put up the “mental hand” in an effort to, as we say, “not go there.” This form of mental control occurs with striking frequency in mental life, especially following disturbing or traumatic experiences (Dougall, Craig, & Baum, 1999), and can certainly lead one to wish that a memory “deletion” device existed. Despite the prevalence of these experiences and their clear clinical importance, experimental psychology has had surprisingly little to say about how this control is accomplished, and what its limits may be.

In this chapter, I have reviewed our approach to this issue, the empirical evidence that supports it, and the relation of these findings to data obtained in research on directed forgetting. Our central claim is that the capacity to control memory rests on the ability to override unwanted memory retrievals, a function that I claim is rooted in a fundamental ability to override prepotent responses. By this view, the moment when we exclude an unwanted memory from awareness is accomplished by the same, or at least highly similar, systems to those that help us to stop a physical action upon demand. By studying this model task, its functional properties, and the brain systems that underlie it, we hope to gain a better understanding of naturally occurring cases of motivated forgetting. In so doing, perhaps we may be able to better the circumstances of individuals for whom memory intrusions are debilitating. Nevertheless, voluntary suppression is unlikely to result in a “spotless mind,” and certainly would not work as quickly as a “memory deletion” device. However, the slower, gradual human solution to forgetting unwanted memories may be a graceful compromise between

the desire to expel what is unpleasant from our lives and the need to retain all experiences to grow as individuals.

Author Notes

The author would like to thank Benjamin Levy and Colin MacLeod for helpful comments on an initial draft of this chapter. Preparation of this chapter was supported in part by grant 5 RO1 MH62126-02 from the National Institute of Mental Health, USA. The author may be contacted by e-mail at mcanders@uoregon.edu.

References

- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, M. C. (2001). Active forgetting: Evidence for functional inhibition as a source of memory failure. *Journal of Aggression, Maltreatment, & Trauma*, 4, 185-210.
- Anderson, M. C. (2003). Rethinking interference theory: Executive control and the mechanisms of forgetting. *Journal of Memory and Language*, 49, 415-445.
- 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, & Cognition*, 20, 1063-1087.
- 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, CA: Academic Press.
- Anderson, M. C., & Green, C. (2001). Suppressing unwanted memories by executive control. *Nature*, 410, 131-134.
- 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. (2004). Neural systems underlying the suppression of unwanted memories. *Science*, 303, 232-235.
- 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.
- Aron, A. R., Fletcher, P. C., Bullmore, E. T., Sahakian, B. J., & Robbins, T. W. (2003). Stop-signal inhibition disrupted by damage to right inferior frontal gyrus in humans. *Nature Neuroscience*, 6, 115-117.

- Aron, A.R., Robbins, T. W., & Poldrack, R. A. (2004). Inhibition and the right inferior frontal cortex. *Trends in Cognitive Sciences*, 8, 170-178.
- Basden, B. H., Basden D. R., & Gargano, G. J. (1993). Directed forgetting in implicit and explicit memory tests: A comparison of methods. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 19, 603-713.
- Bjork, E. L., & Bjork, R. A. (2003). Intentional forgetting can increase, not decrease, residual influences of to-be-forgotten information. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 29, 524-531.
- 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: Lawrence Erlbaum Associates.
- Bjork, E. L., Bjork, R. A., & Glenberg, A. (1973, November). Reinstatement of interference owing to to-be-forgotten items. Paper presented at the 14th Annual Meeting of the Psychonomic Society, St. Louis, MO.
- 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.
- Conway, M. A., & Fthenaki, A. (2003). Disruption of inhibitory control of memory following lesions to the frontal and temporal lobes. *Cortex*, 39, 667-686.
- 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.
- 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 & Dissociation*, 2, 67-82.
- Dougall, A. L., Craig, K.J., & Baum, A. (1999). Assessment of characteristics of intrusive thoughts and their impact on distress among victims of traumatic events. *Psychosomatic Medicine*, 61, 38-48.
- Eldridge, L. L., Knowlton, B.J., Furmanski, C.S., Bookheimer, S. Y., & Engel, S. A. (2000). Remembering episodes: A selective role for the hippocampus during retrieval. *Nature Neuroscience*, 3, 1149-1152.
- Freud, S. (1966). In Strachey, J. (Ed.), *The Standard Edition of the Complete Psychological Works of Sigmund Freud* (pp. 117-128). London: 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.
- Garavan, H., Ross, T. J., Murphy, K., Roche, R. A. P., & Stein, E. A. (2002). Dissociable executive functions in the dynamic control of behavior: Inhibition, error detection, and correction. *NeuroImage*, 17, 1820-1830.
- Geiselman, R. E., Bjork, R. A., & Fishman, D. L. (1983). Disrupted retrieval in directed forgetting: A link with posthypnotic amnesia. *Journal of Experimental Psychology: General*, 112, 58-72.

- Goernert, P. N., Larson, M. E. (1994). The initiation and release of retrieval inhibition. *Journal of General Psychology*, 121, 61-66.
- Golding, J. M., & MacLeod, C. M. (Eds.). (1998). *Intentional forgetting: Interdisciplinary approaches*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Hicks, J. L., & Starns, J. J. (2004). Retrieval-induced forgetting occurs in tests of item recognition. *Psychonomic Bulletin & Review*, 11, 125-131.
- Kimball, D. R., & Bjork, R. A. (2002). Influences of intentional and unintentional forgetting on false memories. *Journal of Experimental Psychology: General*, 131, 116-130.
- Knight, R. T., Staines, W. R., Swick, D., & Chao, L. L. (1999). Prefrontal cortex regulates inhibition and excitation in distributed neural networks. *Acta Psychologica*, 101, 159-178.
- McGeoch, J. A. (1942). *The psychology of human learning: An introduction*. New York: Longmans.
- Menon, V., Adelman, N. E., White, C. D., Glover, G. H., & Reiss, A. L. (2001). Error-related brain activation during a Go/NoGo response inhibition task. *Human Brain Mapping*, 12, 131-143.
- 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.
- Myers, L.B., Brewin, C. R., & Power, M. J. (1998). Repressive coping and the directed forgetting of emotional material. *Journal of Abnormal Psychology*, 107 (1), 141-148.
- Nickerson, R. S. (1984). Retrieval inhibition from part-set cuing: A persisting enigma in memory research. *Memory & Cognition*, 12, 531-552.
- 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.
- 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.
- Rundus, D. (1973). Negative effects of using list items as retrieval cues. *Journal of Verbal Learning and Verbal Behavior*, 12, 43-50.
- Sahakyan, L. & Kelley, C. M. (2002). A contextual change account of the directed forgetting effect. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 28, 1064-1072.
- Sakagami, M., & Niki, H. (1994). Spatial selectivity of go/no-go neurons in the monkey prefrontal cortex. *Experimental Brain Research*, 100, 165-169.
- Shimamura, A. P. (2000). The role of the prefrontal cortex in dynamic filtering. *Psychobiology*, 28, 207-218.
- Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review*, 99, 195-231.
- Veling, H., & van Knippenberg, A. (2004). Remembering can cause inhibition: Retrieval-induced inhibition as cue independent process. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 30, 315-318.

- Zacks, R. T., Radvansky, G., & Hasher, L. (1996). Studies of directed forgetting in older adults. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 3, 231-237.