Theoretical issues in inhibition:

Insights from research on human memory

Michael C. Anderson & Benjamin J. Levy

University of Oregon

Address Correspondence to:

Michael C. Anderson
Department of Psychology
1227 University of Oregon
Eugene, Oregon, 97403-1227
mcanders@uoregon.edu

Theoretical issues in inhibition:

Insights from research on human memory

Several years ago one of the authors of this chapter (BL) married a woman he had been dating for several years. On numerous occasions after the wedding, he mistakenly referred to his new wife as his "girlfriend." During the years of courtship he had become accustomed to referring to her in this manner, so much so that the word just popped to mind when thinking of her. The recent marriage made that label problematic, however. On several occasions when she was not present, he made public references to his "girlfriend," leaving others to assume he had taken up a mistress and was simply very candid about the fact. While amusing at first, the humor of these mistakes was quickly lost on his wife. With effort and attention, however, he was able to override this well-practiced response and refer to her as his "wife." This example illustrates a simple point: often we are victims of an overly effective retrieval system. The years of practice at referring to her as his "girlfriend" made retrieving that label essentially automatic.

While the above example is amusing, oftentimes we encounter reminders to things that we are much more motivated to avoid thinking about, such as a painful breakup, a particularly odious task that we would prefer not to do, or a loved one who has passed away. In these circumstances, we may exert effort to prevent these memories from occupying our thoughts. How are we able to prevent intrusive, unwanted memories from coming to mind? Research from our laboratory suggests that people recruit executive control to override the retrieval of unwanted memories and that this cognitive act induces a lasting suppression of the unwanted memories, making them more difficult to recall later even when we want to return to them (see Anderson, 2003 for a review).

A core theme of this research is that our ability to control unwanted memories is directly analogous to our ability to control our overt behavior, a topic broadly studied in cognitive psychology and in cognitive neuroscience. In fact, the situations described above resemble a classic situation that requires executive control, often referred to as response override (see Fig 1). In response override situations one must stop a prepotent response to a stimulus, either because that response must be withheld or because an alternative, more weakly learned response to that stimulus is desired. The ability to stop prepotent responses is critical to the flexible control of behavior, whether the response is that of a baseball player stopping the swing of his bat when the pitch is a ball, that of a husband avoiding embarrassment, or that of a person preventing an unpleasant memory from coming to mind. Without this ability we would be slaves of our habits and reflexes. Our suggestion, one in keeping with other research on executive control, is that we accomplish this control through inhibition of the prepotent response. When we are presented with reminders to unwanted memories, activation spreads from the cue to the traces stored in memory. If the dominant trace is not currently desired, either because the rememberer wishes to avoid thinking about it or because a more weakly associated trace is sought, then inhibition can be engaged to weaken the dominant memory, enabling us to stop retrieval or to have selective control over what we retrieve. We argue that these inhibitory mechanisms also have a lasting effect, leading to later memory impairments for the avoided memories even when we want to retrieve them later.

-----INSERT FIGURE 1 ABOUT HERE-----

In this chapter we discuss evidence that uniquely supports this inhibitory control perspective in two response override situations in memory: the desire to stop retrieval, and the need to selectively retrieve a memory. Then we describe a theoretical problem in the measurement of inhibition that is almost always ignored in studies of inhibition. This issue, the correlated costs and benefits problem, has extremely important consequences for the ability to adequately test theoretical models of inhibition, and, importantly, for the ability to test inhibitory deficit theories concerning different populations of subjects. We argue that in order to make a strong claim in any study about the presence or absence of inhibition, or about variations in the magnitude of inhibition as a function of condition or population, it is *necessary* to include an independent probe of the impaired items' accessibility. Without this, measurements of inhibition will suffer the correlated costs and benefits problem, precluding principled predictions about how behavioral effects should vary according to inhibitory theories. This problem is not at all unique to memory research, and examples of the problem in research on executive control, visual selective attention, and language processing will be provided. By drawing attention to this issue, we hope to steer the field toward experiments that isolate the involvement of inhibitory mechanisms in the control of memory retrieval and prevent the unnecessary confusion and theoretical controversy in the literature on inhibitory processes.

Selective Retrieval

When recalling a memory, the desired trace is rarely the only memory related to the cues guiding retrieval. Most cues are related to many memories, and very often, non-target memories are more strongly associated to the cue than is the currently desired trace. In this situation, the associated traces compete for access to conscious awareness, necessitating some process to enable selective retrieval. In our framework, selective retrieval represents a paradigmatic case of response override, where one must select a weaker memory in the face of interference from one or more prepotent competitors. If stopping prepotent responses engages inhibition, the same mechanisms might also be engaged to stop prepotent memories from coming to mind, promoting selective retrieval. If so, perhaps inhibition will induce long-lasting memory impairment for the competitors. Thus, the very act of remembering should cause forgetting of related memories.

This prediction has been explored in a procedure we refer to as the retrieval practice paradigm. In a typical version of this paradigm, subjects study lists of category-exemplar words pairs (e.g., Fruits-Orange, Fruits-Banana, Drinks-Scotch) and then they practice retrieving half of the studied items from half of the categories (e.g., “Fruits-Or”), each of which is practiced three times. After a delay, subjects are asked to recall all of the previously studied exemplars. Not surprisingly, practiced items (Orange) are facilitated relative to items from non-practiced categories (Scotch), which act as a baseline measure of how well items are recalled when there is no practice. More interestingly, the non-practiced items from the practiced categories (Banana) are recalled less often than the baseline items (see Fig 2). Thus, retrieving some items during retrieval practice leads to worse memory for related items on the final test. According to the inhibitory control hypothesis, this occurs because these items are inappropriately activated during the retrieval practice phase and then are inhibited to promote successful retrieval of the desired response (Orange). This phenomenon, known as retrieval-induced forgetting (RIF), has now been replicated many times using a broad array of stimuli (for reviews, see Levy & Anderson, 2002, and Anderson, 2003). A brief survey of recent demonstrations of RIF illustrates the striking breadth of this phenomenon: when bilinguals retrieve the non-native word for a concept, the native language phonology of that word is inhibited (Levy, McVeigh, Marful, & Anderson, submitted), retrieving multiplication facts about a particular number (e.g., 7 times 7 equals 49) impairs the ability to report unpracticed multiplications with that number (e.g., that 7 times 8 equals 56) (Phenix & Campbell, 2004), and retrieving individuating characteristics of a person from a stereotyped group causes inhibition of the target's stereotypic traits (Dunn & Spellman, 2003). More relevant to theoretical explanations of the effect, recent research demonstrates that while most RIF studies induce competition by episodic exposure to the competitors, semantically related items can be inhibited even when not studied (Johnson & Anderson, 2004; Starns & Hicks, 2004). Thus, a wealth of experimental results suggests that RIF is a general phenomenon that results in impairment whenever unwanted items intrude during retrieval.

-----INSERT FIGURE 2 ABOUT HERE-----

One of the goals of this volume, however, is to determine whether inhibition is actually involved in "inhibitory" effects such as RIF. The basic finding of RIF described above is compatible with several non-inhibitory mechanisms and by itself is not compelling evidence that inhibition is involved. For example, the practiced items may be so strengthened by retrieval practice that they interfere during the final test, effectively blocking the subject from coming up with the correct response. This type of retrieval competition has a long history in formal models of memory retrieval (e.g., Raaijmakers & Shiffrin, 1981; Anderson, 1983), where the probability of recalling an item is predicted by the relative strength of the association between the cue and the target compared to the strength of the association between the cue and all the competitors. Other non-inhibitory mechanisms can also explain RIF as well. For example, subjects may simply unlearn the connection between “Fruits” and “Banana” or practicing some exemplars may bias the representation of that category toward the practiced items (retrieving “Orange” might bias the cue “Fruits” towards “Citrus Fruits”). According to all of these non-inhibitory explanations, one does not need to claim that any change is occurring to the item itself (see Anderson & Bjork, 1994, for a review of these non-inhibitory sources of impairment). However, this is a core claim of the inhibitory control perspective: the actual unwanted item itself is being made less accessible.

Several properties of RIF uniquely support the involvement of inhibition. First, the inhibitory control perspective (Anderson, 2003) makes the unique claim that retrieval-induced forgetting should be observed regardless of which cue is used to test the memory. In other words, forgetting should be *cue independent* and should generalize to novel cues in the test phase, rather than being specific to those used to perform retrieval practice (Anderson & Spellman, 1995). For example, recall of “Banana” should be impaired not only when it is tested with the studied category (“Fruits”), but also when it is tested with a novel, independent retrieval cue (Monkey-B). Such cue-independent forgetting is difficult for non-inhibitory mechanisms to explain since they predict that impairment should be specific to the cues used during retrieval practice. For example, associative blocking cannot explain this type of impairment because the new independent cue (Monkey) is unrelated to the strengthened exemplars (e.g., Orange), thus there is no reason why they should block retrieval of “Banana”. Cue-independent forgetting has now been observed many times (e.g., Anderson & Spellman, 1995; Radvansky, 1999; Anderson & Bell, 2001, Anderson, Green, & McCulloch, 2000; Camp, Pecher, & Schmidt, 2005; MacLeod & Saunders, 2005), suggesting that the competing item itself is inhibited. Other researchers have expanded on the notion of *cue independence* by arguing that if competitors are truly inhibited, then memory impairments should also occur on other types of memory tests besides just recall. RIF has now been found on both recognition tests (Hicks & Starns, 2004; Starns & Hicks, 2004) and an implicit lexical decision test (Veling & van Knippenberg, 2004). The use of independent probes is the best tool that we have at present for distinguishing between interference and inhibition as mechanisms in long-term memory. Despite this fact, many researchers design studies of inhibition in long-term memory without using independent probes. The problem with this approach is that one cannot isolate the role of inhibition in producing the memory impairment without a way of excluding contributions of non-inhibitory factors. Lack of attention to this fact has led, we argue, to misleading conclusions about the nature of retrieval-induced forgetting, and how it varies in different populations of subjects. We will return to this concept in much greater detail later in this chapter.

Second, a hallmark of associative interference accounts of forgetting is that memory impairment should occur whenever competitors are strengthened. Several findings, however, have shown that strengthening practiced items actually does not always lead to impaired recall of the competitors. For example, repeated study exposures to the practiced items results in strengthening comparable to that produced by retrieval practice, yet this type of strengthening produces no forgetting of the competing items, provided that output interference at test is controlled (Anderson, Bjork, & Bjork, 2000; Ciranni & Shimamura, 1999; Bäuml, 1996, 1997, 2002). If strengthening can occur without producing impairment of the related items, it suggests that retrieval practice, not strengthening, is what causes the impairment. Thus, RIF appears to be both *recall specific* and *strength-independent.*

Lastly, the amount of impairment depends on the extent to which the competitors interfere, a property we have previously referred to as *interference-dependence*. For example, Anderson, Bjork, and Bjork (1994) found that more RIF is observed when the competitors are high frequency members of the category (e.g., Orange) than when they are uncommon examples (kiwi). Similarly, doing retrieval practice on the meanings of asymmetric homographs results in significant impairment for the dominant meaning, but not for the subordinate meaning (Shivde & Anderson, 2001). This also poses problems for non-inhibitory accounts of RIF. For example, according to the associative blocking account there is no reason to expect that the relative frequency of the unpracticed item should influence the degree to which the practiced item blocks retrieval during the final test. Each of these properties of RIF strongly implicates inhibition and provides serious challenges for any non-inhibitory explanation of RIF. Taken together, these results support the inhibitory control perspective that selective retrieval is a special case of response override that results in lasting inhibition of the avoided memories.

Stopping Retrieval

The need to control behavior is not only limited to selecting a non-dominant response, it is also sometimes necessary to stop a behavior from occurring. For example, when a baseball player makes the last minute decision not to swing at a ball that is out of the strike zone, he or she is choosing to override a prepared response to swing. Sometimes our goal as rememberers is to do something remarkably similar: we are not trying to select a competing memory so much as we try to simply stop retrieval itself.

When presented with a reminder of something upsetting, for example a photograph of a loved one who has recently died, we might desire to put those thoughts out of mind. Other times, we may simply want to stop a thought from popping into mind, because we are trying to stay focused on a different thought. In fact, this type of control of internal thought is crucial for preventing our minds from wandering. Can inhibitory mechanisms be engaged to serve these goals?

Evidence for the ability to stop memory retrieval comes from the "think/no-think" (TNT) paradigm, developed by Anderson & Green (2001). In a typical TNT study, participants learn cue-target pairs (e.g., ordeal-roach). Then subjects are presented with the cue they studied earlier (e.g., ordeal) and are either asked to think of the associated target memory (e.g., roach) or to prevent that associated target from coming to mind. Subjects are instructed that the target memory is not supposed to come to mind even for a moment. After seeing these cues as many as 16 times, subjects are asked on a later memory test to recall all of the previously studied target memories. When provided with the same cue that they studied earlier (ordeal) and asked to provide the correct target memory (roach), subjects have more difficulty if they had previously avoided thinking about that memory than if they had encountered no reminders to it during the TNT phase (see Fig 3). This is counterintuitive, because one would expect that repeated reminders to an item should make it more accessible, not less accessible.

-----INSERT FIGURE 3 ABOUT HERE-----

Just as was the case in RIF, the basic TNT suppression effect is compatible with non-inhibitory accounts. For example, subjects might generate alternative thoughts during the TNT phase for each of the cues that require suppression (e.g., some "ordeal" they experienced personally). Then on the final test these diversionary thoughts might come to mind, blocking retrieval of the target item. To rule out this possibility, Anderson and Green (2001) demonstrated memory impairment for the avoided items using novel, extralist retrieval cues (e.g., "insect-r______" for "roach") as independent probes. Associative interference from diversionary thoughts cannot explain independent probe impairment, because the diversionary thoughts were generated to the originally studied cue. This finding implicates inhibition as the mechanism that excludes unwanted memories and that decreases their accessibility. Anderson and Green also ruled out alternative explanations by replicating the memory impairment even when subjects were paid for correct responses or were misled about the expectations of the experimenters (i.e., that attempting to suppress an idea would ironically make it more accessible). In their final experiment, they showed that simply asking subjects to avoid saying the response word eliminates the suppression effect, indicating that the attempt to prevent the unwanted memory from coming to mind is critical to producing the impairment. Taken together, these results indicate that the memory impairment arises from an inhibitory control mechanism that prevents the unwanted memory from entering awareness.

Think-no think suppression has now been replicated several times (Anderson et al, 2003, Hertel & Gerstle, 2003 – non-dysphoric subjects on positive word pairs; Hertel & Calcaterra, 2004 – "aided" condition; Joorman, Hertel, Brozovich, & Gotlib, 2005; Wessel, Wetzels, Jelicic, & Merckelbach, 2005; Depue, Banich, & Curran, in press; Anderson, Reinholz, Kuhl, & Mayr, in preparation; Hota & Kawaguchi, 2006; Kawaguchi & Hota, 2006, Anderson & Kuhl, in preparation; Bell & Anderson, in preparation; although, see Bulevich, Roediger, Balota, & Butler, in press, for a failure to find below baseline impairment). While Anderson and Green (2001) used neutral pairs, several recent studies have investigated whether the valence of the unwanted thought influences suppression success. Hertel and Gerstle (2003) found that non-depressed subjects were able to suppress positive adjective-noun pairs, but were not able to suppress negative pairings (depressed undergraduates could not suppress either type of stimuli). In contrast to this failure to suppress negative stimuli, three more recent studies have found that suppression of negative stimuli leads to comparable (actually non-significantly increased) inhibition relative to either neutral stimuli (Depue et al, in press; Anderson et al, in prep) or positive stimuli (Joorman et al, 2005). It is unclear what produced the different pattern in Hertel and Gerstle's study, but the simplest conclusion seems to be that the valence of the unwanted memory itself may not be the primary determinant of the amount of suppression.

The memory impairment observed in TNT studies suggests that inhibitory control mechanisms may be recruited to prevent unwanted memories from coming to mind. This finding has obvious implications for situations where people wish to avoid persistent, intrusive thoughts. One approach might be for people to simply avoid reminders to those memories. However, this is often not practical. This research suggests that when placed in an environment with unavoidable reminders, people may be able to control their memories and prevent the unwanted traces from coming to mind. By this view, repeatedly avoiding memories in naturalistic settings may cause long-lasting impairments at recalling those memories. Thus, the TNT paradigm may be a useful laboratory model of the voluntary form of repression (suppression) proposed by Freud (1966). If so, the

TNT paradigm, along with other procedures such as directed forgetting, can be used as a tool for exploring clinical issues related to motivated forgetting.

Neural substrates of stopping retrieval

Earlier we outlined how stopping retrieval is analogous to stopping a motor response. A key difference between these two situations is that motor suppression is observable with the naked eye (e.g., you can see when a baseball player "checks" his swing), but there is no outwardly observable sign when someone chooses to stop a declarative memory. Despite the surface similarities between stopping actions and stopping thoughts, they could be accomplished in entirely different ways. One way of testing our idea about a fundamental similarity between these situations is by investigating whether these processes are produced by the same brain mechanisms. If stopping retrieval is really similar to stopping a motor response, then a common underlying neural network should be involved in accomplishing both types of stopping.

Studies of motor response override have shown that a network of control-related regions, including the lateral prefrontal cortex, anterior cingulate cortex, lateral premotor cortex, and intraparietal sulcus are recruited to stop motor responses (e.g., Garavan et al., 2002; Menon et al., 2001). Thus, suppressing unwanted memories might also engage these regions. In addition, there should also be some regions unique to memory control: those regions that are the target of suppression. Given the role of the hippocampus in conscious recollection of declarative memories (e.g., Squire, 1992; Eldridge et al, 2003) and the goal of suppressing conscious recollection in the think/no-think task, the hippocampus seems a likely region to target. Based on the foregoing analysis, Anderson et al (2004) hypothesized that control-related regions (particularly lateral prefrontal regions) should be involved in disengaging hippocampal processes to prevent conscious recollection of the unwanted memories.

Anderson et al (2004) addressed this hypothesis by using fMRI to identify the brain regions that support intentional memory suppression. They had subjects participate in a procedure similar to the one described earlier and they scanned subjects while they performed the TNT task. They found that their subjects recalled significantly fewer suppression than baseline words, replicating the earlier behavioral work. To investigate the hypotheses concerning neural regions, they compared the neural response to cue words that required them to think of the response word to ones that required them to suppress the corresponding response word (see Fig 4). Perhaps the most striking finding was that there was significantly more activation in control-related regions during suppression trials than during the respond trials. This demonstrates that stopping retrieval is not simply a failure to engage retrieval processes. Instead, that these regions showed heightened activation during suppression trials suggests that subjects must actively engage processes to prevent the unwanted memories from coming to mind. As predicted, the observed regions of increased activity overlapped highly with control-related regions involved in stopping motor responses, confirming the idea that common brain regions control stopping both unwanted memories and unwanted actions.

-----INSERT FIGURE 4 ABOUT HERE-----

In addition to these control regions, Anderson et al (2004) observed the predicted reduction in hippocampal activity bilaterally on “no think” trials relative to “think” trials. While this difference could be due to increased activity during the “think” condition, it is also consistent with the idea that the hippocampus is down-regulated on suppression trials. Supporting the latter explanation, the degree of hippocampal activity was related to behavioral memory inhibition (see Anderson et al, 2004, for a description of this relationship). The fact that hippocampal activity is correlated with behavioral suppression is not consistent with the idea that the difference between hippocampal activity on think and no think trials is *entirely* due to heightened recollection in the think condition. Instead, it suggests that subjects can strategically down-regulate the hippocampus to prevent conscious recollection.

The results of this imaging study suggest that subjects can prevent unwanted memories from coming to mind by the same neural mechanisms that are recruited to stop motor actions. Instead of targeting a motor response, these regions are recruited to suppress declarative memory representations. So, despite the difficulty involved with "observing" the stopping of memory retrieval, we now have both behavioral and neuroanatomical evidence to support the idea that people can prevent unwanted memories from coming to mind. The findings thus help to specify a model of how motivated forgetting might occur.

The Correlated Costs and Benefits Problem: A Central Theoretical Issue for Research on Inhibitory Control

Although the foregoing studies have isolated the role of inhibition in memory retrieval, they have also revealed significant theoretical issues in how inhibition should be measured. In the last part of this chapter, we discuss one of the most important issues that has wide implications for investigators interested in inhibition: the *correlated costs and benefits problem*. We argue that failure to attend to this issue has contributed to significant confusion in the literature, especially in work on individual differences in inhibitory function. Such confusion has hindered theoretical development, and generated doubt about inhibition phenomena—doubt that may underlie some of the skepticism evinced by opponents to inhibitory theories. We offer a solution that has proven useful in research on memory inhibition, and that could be adapted to the study of inhibition in other cognitive domains.

The Nature of the Problem.

All investigators would agree that whatever cognitive mechanisms are proposed to explain a phenomenon ought to operate consistently: If a mechanism is engaged under one set of circumstances, it ought to be engaged in future situations that have similar circumstances. Surprisingly, this seemingly uncontroversial premise is routinely ignored in studies of inhibition. Studies that use some procedure to induce inhibition at time T, and later measure the behavioral aftereffects of that inhibition at time T+1, very frequently fail to consider the involvement of inhibition at time T+1, even though such mechanisms often contribute to the measurement process itself.

Most investigators agree that if inhibition degrades a representation at time T, performance may be hindered on that item at time T+1; however, when that aftereffect is measured at T+1, the inhibited item itself becomes the target of processing and those same inhibitory mechanisms may assist in retrieving/processing it, by suppressing interference from *other* items that may be competing during that test. When the T+1 measurement provides such an opportunity to profitably engage inhibition, the correlated costs and benefits problem arises. *The correlated costs and benefits problem refers to the theoretically predicted difficulty in quantifying the amount of inhibition that has taken place at time T by measuring inhibition using an assessment at time T+1 that itself* *requires the same inhibitory process.* Under these circumstances, the same inhibitory process that at time T induces behavioral aftereffects (usually costs to later performance) can be brought to bear at time T+1 to mitigate competition (i.e., introducing benefits at test). Because these hypothesized costs and benefits stem from the same inhibition process, increments in inhibitory control (as generated by different populations or experimental conditions) should induce costs that are, in general, offset or balanced by decrements in non-inhibitory sources of interference at time T+1. Thus, the mapping between the behavioral performance decrement and inhibitory control function becomes impossible to determine, compromising the measurement of inhibition.

To see this, consider Figure 5a, which uses a retrieval task to illustrate the point. In each case, a cue is associated to two items in memory, and subjects' task is to retrieve one of those items, either item A in Situation A or item B in Situation B. According to the inhibitory control perspective, in Situation A if the cue activates both traces A and B, they should compete for access to awareness; thus if item B interferes with the retrieval of item A, inhibition should suppress item B, facilitating the retrieval of item A. The lingering effects of inhibition on item B should make it harder to recall on later occasions. In Situation B, we have the same circumstances--only the particular item that is to be recalled is item B rather than item A, leading item A to be the target of inhibition. No inhibition theorist would disagree that inhibitory control would be involved in both of these equivalent retrieval situations (assuming that interference is present in both).

-----INSERT FIGURE 5 ABOUT HERE-----

Although this reasoning seems obvious when laid out in this fashion, it has not been as obvious to investigators studying the role of inhibition in memory retrieval using retrieval-induced forgetting. To see this, consider Figure 5b, which illustrates different phases of the retrieval-practice paradigm. In each panel we have a cue related to two targets in memory, as in Figure 5a. In the retrieval practice phase, subjects practice retrieving Orange, suppressing the competitor Banana and making it less accessible. On the final test we have the same items, only now the goal is to retrieve the former competitor, which is now the target (e.g., Banana). Clearly, consistency demands that the same inhibitory control mechanisms used to suppress Banana in the retrieval practice phase ought to now be engaged to suppress Orange, which has become the competitor. Indeed, inhibitory control should be more necessary during the test both because Banana is inhibited (and harder to retrieve) and because Orange is stronger from retrieval practice. Thus, the same inhibitory mechanisms that impaired Banana during retrieval practice (time T), inducing a cost, should now during the final test (time T+1), aid in the retrieval of Banana, by overcoming interference from the previously practiced item (Orange). Importantly, the magnitude of both the cost and the benefit should be correlated, because they arise from the same underlying mechanism: inhibition.

To see the impact of the correlated costs and benefits problem, consider the example given in Figure 6, using retrieval-induced forgetting as a model case. Suppose that inhibition contributes to retrieval-induced forgetting (for example), and that individuals vary in their inhibition ability. Suppose further that we knew, for each person, how much of their retrieval-induced forgetting was produced by suppression during retrieval practice, and how much was attributable to blocking on the final test (i.e., inability to combat the interfering effects of the strengthened practiced item). What would the predicted relationship between these components and inhibitory control ability look like? How would they sum together to produce the aggregate retrieval-induced forgetting effect? In Figure 6, the x-axis plots inhibitory control ability, and the y-axis plots the amount of retrieval-induced forgetting attributable to either suppression (dark line), blocking (dotted line), or to both factors (dashed line). As illustrated in Figure 6, the suppression component should increase with inhibitory ability. So, for example, a patient with damage to the prefrontal cortex might be positioned on the left side of this figure, and be unable to suppress competing items during retrieval practice, whereas a high functioning college-aged student might be closer to the right side. If retrieval-induced forgetting were solely produced by suppression induced during retrieval-practice, then the effect should clearly increase with increasing inhibitory control ability.

-----INSERT FIGURE 6 ABOUT HERE-----

Unfortunately, the correlated costs and benefits problem implies that the same inhibitory control ability that suppresses the competitor during retrieval practice (time T) ought to assist in suppressing the practiced item during the final test (time T+1), reducing the blocking contribution to the phenomenon. Thus, as inhibitory control ability increases from the left to the right, the component of retrieval-induced forgetting attributable to blocking should diminish, reflecting subjects' growing ability to combat that interference on the final test. Thus, on the left side of this figure, a hypothetical frontal patient, even if they have absolutely no ability to suppress the competitor during retrieval-practice, might, for the same reasons, be completely unable to suppress interference from the practiced items on the final test. To the extent that practiced items are strengthened, the frontal patient should experience exaggerated blocking from those strengthened items. Ignoring this fact would be especially curious, since perseverative behavior in memory tasks is largely what motivated the original inhibitory deficit hypothesis of frontal lobe dysfunction. On the right side of the figure, by contrast, the hypothetical college student, while they may have effectively suppressed the competitor during practice, will suffer little associative blocking on the final test. When the joint contributions of the costs (persisting suppression) and benefits (reduced blocking at T+1) are summed (dark line), one can see that inhibitory control theories do not predict a clear relationship between inhibitory control ability and the size of the retrieval-induced forgetting effect when subjects are tested with the same cue used to perform retrieval practice. Essentially, the costs and benefits should trade off, yielding effects that appear the same behaviorally, but that are generated by different underlying chains of events.

Examples of Failures to Consider This Problem, and Their Consequences.

Although this analysis may seem straightforward, several studies that have used retrieval-induced forgetting to test inhibitory deficit hypotheses have not considered the role of suppression on the final test. These studies illustrate the potential for theoretical confusion introduced by this problem. In each study, the same-probe test of retrieval-induced forgetting was employed, in which the potential contributions of blocking on the final test leads to a significant correlated costs and benefits problem. Consider, for example, studies by Moulin et al. (2002) and Conway and Fthenaki (2003), both of which measured retrieval-induced forgetting with the standard category cued recall final test procedure (Figures 7a and 7b). These investigators contrasted a patient population thought to have a deficit in inhibitory control (either Alzheimer's patients, or Frontal-

Lobe Damaged patients) against a control population (normal older adults, or patients with other brain damage). In both studies, retrieval-induced forgetting was observed in the deficit and control populations; in fact, in some cases, more impairment was found in the population with the supposed deficit than in the population with intact inhibitory function. If retrieval-induced forgetting is a pure measure of inhibition, then either these populations do not have deficits in inhibitory control, or the inhibition measured by retrieval-induced forgetting is different than that supported by the prefrontal cortex. Indeed, these and other authors have used evidence such as this to argue that retrieval-induced forgetting reflects an automatic form of inhibition, not mediated by prefrontal cortex, and not related to attentional control (Zellner & Bäuml, 2005). Yet, if one invokes the inhibition process consistently when it is required, there should be no clear relationship between retrieval induced forgetting and the deficit status of these populations, as measured by the same probe test. Put simply, the Alzheimer's disease and frontal lobe damaged patients, even if they were completely lacking in inhibitory control, would be expected to show significant forgetting, by virtue of their compromised ability to resolve interference at the time of the final test. This makes their retrieval-induced forgetting effect appear similar to control subjects with intact inhibitory function, despite their impaired recall being produced by an entirely different cause.

-----INSERT FIGURE 7 ABOUT HERE-----

The potential result of failing to consider the correlated costs and benefits problem, as illustrated here, is to lead to theoretically significant conclusions about the nature of retrieval-induced forgetting, or, alternatively, about the populations being studied—conclusions that, based on a more thorough consideration of correlated costs and benefits analysis, are not warranted and *potentially* very misleading. Evidence of this potential to mislead is provided next, along with a solution to the problem.

The Correlated Costs and Benefits Problem: A Solution with Examples.

The key to solving the correlated costs and benefits problem is to redesign the final assessment of inhibition to eliminate or reduce the *benefits* of inhibition. Here again, this can be illustrated with memory inhibition, although the principle applies generally for all paradigms in which inhibition is measured. In retrieval-induced forgetting, this approach would require devising a final T+1 assessment in which the heightened interference from practiced items could be minimized, allowing the test to better reflect suppression that took place in the retrieval practice phase (time T). Fortunately, this test already exists: the independent probe method. In the independent probe method, an item (e.g., Fruit Banana) that has been inhibited by retrieval practice of a competing target (e.g., Fruit Orange) is tested with a distinct cue (e.g., Monkey B___) unrelated to the retrieval practice cue (Fruit) and practice target (Orange). Testing the inhibited item with a novel cue circumvents associative interference that might otherwise arise from the stronger practiced items, permitting a measure of inhibition at time T+1 that is not contaminated by blocking. If this test truly minimizes blocking, then there should be less need to engage inhibitory control on the final test, reducing the benefits of that process. Thus, with the independent probe method, the measured retrieval-induced forgetting should better reflect inhibition caused by retrieval practice (time T), making performance on this test more sensitive to individual differences in inhibition capacity than performance on the same probe test, in which costs and benefits are both expressed.

The independent probe method should also increase sensitivity to individual differences in inhibition when used in the think/no-think paradigm. Although subjects in this paradigm do not explicitly strengthen a competitor during the think/no-think phase (unlike in the retrieval practice paradigm), subjects may generate diversionary thoughts that become associated to the cue word, introducing a blocking component to the same-probe effect (see Hertel & Calcaterra, 2005 for this argument). If so, independent probe impairment should better track individual differences in inhibitory function. This predicted finding has now been obtained in three series of experiments using the think/no-think procedure. These studies examined the size of the forgetting effect as a function of working memory capacity (Bell & Anderson, in preparation), age (Anderson, Reinholz, Mayr, & Kuhl, in preparation), and previous trauma history (Anderson & Kuhl, in preparation), all factors that might moderate how well someone can suppress distracting memories. In each series, inhibition was assessed with both the same test cue that was initially used to induce inhibition (Same Probe condition), as well as with a novel, independent probe (Independent Probe condition), permitting greater isolation of the persisting inhibition component. If working memory, age, and traumatic experience modulate inhibitory control ability, and if the correlated costs and benefits analysis is correct, differences in inhibition should be revealed more on the independent probe test than on the same probe test.

-----INSERT FIGURE 8 ABOUT HERE-----

Figure 8 presents the results from the same probe test and the independent probe tests as a function of working memory capacity, age, and level of prior traumatic experience. The data presented here have been simplified by collapsing over all the individual experiments within each series (one experiment for working memory, 2 for the aging series, and 2 for the trauma series), though the pattern is replicated in each individual study. These panels reveal a pattern that is highly consistent with the correlated costs and benefits problem. In each case, the amount of impairment on the same-probe test is, at best, weakly consistent with an inhibitory deficit: lower working memory capacity subjects show moderately less impairment than higher working memory capacity subjects; older adults show numerically less impairment than younger adults; and persons who experienced fewer traumatic experiences in their lives showed numerically less impairment than persons who experienced more trauma and who might be more skilled at inhibition. However, in all cases, the impairment does not vary reliably with subject group. In contrast, in all three series, the independent probe test revealed significant differences in inhibition as a function of subject group. In fact, the differences in inhibition as a function of age, working memory, and trauma, are quite striking on the independent probe measure, and the three-way interaction of inhibition, subject group, and test-type is significant in each case. Thus, whether an inhibitory deficit was observed in all three studies varied reliably as a function of test type, exactly as predicted by the correlated costs and benefits analysis.

The findings from these three series of studies illustrate the dangers of the correlated costs and benefits problem. Clearly, if we had based our conclusions about the relative inhibitory abilities of these different populations on performance on the same probe test, as did Moulin et al. (2002) and Conway and Fthenaki (2003), we might have concluded that these groups did not differ in inhibitory function, and that inhibition is unrelated to working memory capacity or executive control. Yet we know from our neuroimaging work (Anderson et al., 2004) that the think/no-think task engages frontal cortical regions widely known to be involved in executive function. The independent probe measure, by virtue of reducing the potential benefits of inhibition at test, provided much more striking evidence for variation in inhibition, in line with prior theoretical expectations. Thus, the independent probe method provides an important solution to the correlated costs and benefits problem that can prevent theoretical confusion.

Generality of the Problem

Although we have used retrieval-induced forgetting and the think/no-think paradigm to illustrate the correlated costs and benefits problem, this problem applies broadly to a variety of paradigms in memory, attention, and language. Here, we provide examples of the problem in each of these settings, so that investigators in those areas might better recognize how this problem might apply to their research.

Episodic and Semantic Memory Paradigms. Aside from retrieval-induced forgetting, several memory paradigms have the characteristic whereby strengthening some associates of a cue impairs later recall of other associates of the cue, such as retroactive and proactive interference in the A-B, A-C design, output interference for items associated to a common experimental or contextual cue, and part-set cuing impairment. To the extent that impairment is measured by a shared cue at test (as it always is), it may reflect the joint costs and benefits of inhibition. Semantic fluency tasks, in which a subject must produce as many exemplars of a category as they can within one minute, also suffer from this problem: as one generates new exemplars, other, yet-to-be-listed exemplars in semantic memory ought to be suppressed, yielding a cost to total fluency performance; on the other hand, that same inhibition process should reduce interference from already-generated exemplars, reducing perseveration, yielding a benefit to performance. Ironically, total fluency performance has often been taken as a measure of executive control ability, and in particular inhibitory control ability, when in fact no straightforward relation between total fluency and inhibitory function is predicted by inhibitory control theories.

Executive Control Paradigms. The correlated costs and benefits problem is likely to arise in executive control tasks thought to be tailored to measure inhibitory effects. Consider, for instance, the phenomenon of backward inhibition in task set switching. In the switching paradigm, subjects are presented with a stimulus, upon which they must perform a task as rapidly as possible (e.g., adding a pair of visually presented numbers). Other trials might cue subjects to perform a different task (e.g., subtraction, multiplication) on the same class of stimuli. Of interest is the speed with which subjects perform a target task (e.g., addition) as a function of whether they had performed that same task (addition) or a different task (multiplication) on the previous trial. It is generally found that people are slower to perform the target task if they had just performed a different task, reflecting the time that it takes for people to switch their task set. It has been proposed that on switch trials, subjects engage inhibitory mechanisms to suppress the preceding task set (e.g., multiplication) so that a new set may be selected. For instance, Mayr and Keele (2000) found that if subjects performed a task (task A, e.g., multiplication) on a stimulus on one block of several trials (block 1), and then switched to a second task (task B, e.g., addition) in the next block (block 2), that switching back to the switched-out-of task (e.g., addition) in the next block (block 3) took significantly longer than if subjects had to switch to a third task set on block 3 (task

C, e.g., subtraction). Thus, having to switch back to a previously rejected set is significantly harder than switching to a new set, suggesting some additional cost associated with having just recently rejected that set. Mayr and Keele (2000) interpreted this as evidence for inhibition in task set switching, and termed the phenomenon backward inhibition.

Backward inhibition has been used to study individual differences in inhibitory function in older adults (Mayr, 2001) and in frontal lobe damaged patients (Mayr et al, in press), on the assumption that backward inhibition provides a clean measure of the effects of inhibition. However, backward inhibition is subject to the correlated costs and benefits problem. On the one hand, the better one's inhibitory control, the slower subjects should be to respond to the recently inhibited task set, because the effects of prior inhibition at time T should persist to T+1. On the other hand, when inhibition is measured on block 3 (T+1), subjects must also suppress interference from the recently engaged task set from block 2, and retrieve the suppressed task set, particularly given that the task set is evoked in part by a common stimulus (numbers). Thus, inhibitory control benefits the subject, helping them to combat interference from the most recent set. Although it might seem that the experimental (task A, followed by B, then A) and control conditions (task A, followed by B, then C) are matched for the need to suppress task B during the final test, this is not the case. In the experimental condition, the strength of task B, relative to the inhibited set A should be greater than the strength of task B, relative to the new task set C in the control condition, making inhibitory control more beneficial in the experimental condition. Thus, inhibition induces both costs and benefits that, when taken together, make it quite difficult to predict a clear relationship between inhibitory function and the backward inhibition effect.

Visual Selective Attention Paradigms. Very similar arguments can be made about the negative priming phenomenon (see Neill & Valdes, 1996; Milliken et al, 1998; Tipper 2001 for reviews) in selective attention. In negative priming, the inhibition of a target on a preceding trial (time T) is typically measured on a probe trial (time T+1) that itself requires subjects to resolve competition. As in backward inhibition, this provides opportunity for inhibition to yield a greater measurement-epoch benefit in the experimental condition, in which the competition is between the suppressed target and a non-suppressed distractor, compared to the control condition, in which the competition is between a non-suppressed target and a non-suppressed distractor. Here, again, the putative costs of inhibitory control are mixed with their benefits, yielding no clear predictions about how negative priming should vary as a function of inhibitory ability.

Language Processing Paradigms. The correlated costs and benefits problem also arises in language tasks that have been used to assess inhibition. Consider, for example, the role of inhibition in lexical ambiguity resolution (see Gorfein, 2001 for a collection of reviews). During sentence processing, interpreting a word that has multiple meanings requires that its contextually appropriate sense be accessed. It is generally thought that competition between the word's multiple meanings must be resolved in some fashion, and this function has been attributed to inhibition by some investigators (Gernsbacher & Faust, 1991; Simpson & Kang, 1994). In one paradigm used to study this issue, subjects view trials composed of a briefly presented prime word, followed by a probe letter string, on which subjects must make a lexical decision. When the prime is a homograph (e.g., Bank), reaction time to a related probe (e.g., River) is speeded, consistent with the idea that the prime accessed the sense of the homograph indicated by the probe. More interestingly, when the following trial re-presents the same prime (Bank), but followed by a probe related to the previously rejected sense (e.g., Money), probe lexical decision times are slower than if the first prime-probe trial had presented an entirely unrelated homograph (e.g., Arm-Shoulder). The slowed response to the previously rejected sense has been taken as evidence that the meaning (Money) was inhibited during the first prime-probe trial, the persisting effects of which was measured on the 2nd prime probe trial.

Unfortunately, however, the assessment of inhibition in this paradigm mixes the lingering effects of inhibition on the prime trial (time T) with the need to overcome interference on the probe trial (time T+1), creating a correlated costs and benefits problem. When the probe trial re-presents the homograph cue word, it is very likely that the most recently selected meaning from the preceding trial will grow very active in response, creating a situation analogous to associative blocking in retrieval-induced forgetting. Thus, the paradigm resembles that same probe test previously discussed. Because of this, processing of the probe on the second trial not only requires subjects to overcome the inhibition of the previously rejected sense, but also combat the heightened interference from the competing sense. Thus, inhibition induces both a cost and a benefit, making it very difficult to make principled predictions about how the impairment measured in this paradigm should be related to inhibitory ability. Similar problems arise in numerous other language processing tasks thought to involve inhibition, including anaphoric reference (Gernsbacher, 1989), and metaphor comprehension (Gernsbacher et al, 2001). Thus, if one's goal is to use such behavioral effects to examine individual differences in inhibition, or even to study the theoretical conditions under which inhibition occurs, one needs a measure that does not conflate these sources of impairment (see Shivde & Anderson, 2001, and Johnson & Anderson, 2004 for examples of the independent probe method in language processing).

Summary.

To the extent that inhibitory control has been assessed with measures that suffer from the correlated costs and benefits problem, as we have argued, the literature should be plagued by inconsistencies in the support of inhibitory deficit theories, generating very reasonable doubt about the utility of these theories. We argue that such variability is not necessarily a sign of weakness or correctness of the theories themselves, but rather a sign of the poorly considered measures that have been used to assess the effects of inhibition. As a solution, we argue for the importance of developing testing methods for measuring inhibition that minimize the potential benefit of inhibition during the test epoch itself. By reducing the benefits of inhibition, variation in inhibitory function can be measured more cleanly. Based on several studies, we argue that the independent probe method provides one such example that has proven to be sensitive as a means of testing individual differences in this function. We believe that the logic underlying this method can be adapted to any domain in which inhibition might be studied.

Concluding Remarks.

The purpose of this volume is to review perspectives on the role of inhibition in cognition. On the one hand, many investigators have presupposed inhibition as part of a broader theory of cognition and, in particular, broader theories of individual differences in cognition. The existence of inhibition seems both plausible and necessary: plausible, because many behavioral effects in attention, memory, and language appear like they might arise from inhibition, and necessary because some process for limiting the influence of distracting representations seems essential---a proposal supported by the widespread involvement of inhibition in the nervous system, and by computational analyses indicating that inhibition is necessary for stability in neural networks (e.g., Easton & Gordon, 1984). On the other hand, skeptics of inhibition correctly emphasize that the mere presence of performance decrements that look like inhibition does not by itself require such processes. Indeed, many behavioral phenomena labeled as "inhibition" may be adequately explained without proposing an additional inhibition process. If so, parsimony dictates that the simpler theory be preferred, though it remains a point of debate whether such alternative non-inhibitory theories are always simpler.

Our approach has been to develop clear standards by which to judge whether a behavioral deficit constitutes inhibition. We have done this in the context of work on memory retrieval. In this chapter, we reviewed this work, which establishes functional properties of memory inhibition that uniquely favor the role of inhibitory mechanisms in memory retrieval, and that suggest that non-inhibitory mechanisms are not sufficient to explain these effects: cue-independence, interference dependence, retrieval-specificity, and strength independence. Each of these properties contributes an important piece of the argument against non-inhibitory alternatives, and collectively, they lend strong support to the view that inhibition overrides prepotent responses in memory. Thus, one of the key imperatives advocated by reasonable skeptics of inhibition---to go beyond the mere equating of a performance decrement with inhibition---has been programmatically addressed in the context of memory retrieval over the last decade. This does not imply that some non-inhibitory mechanisms do not also contribute to apparent inhibitory effects in memory, only that inhibition clearly contributes to producing such effects.

In fact, the potential contributions of non-inhibitory sources of impairment to apparent inhibitory phenomena motivated our discussion of what we consider an extremely important problem in the measurement of inhibition: the correlated costs and benefits problem. Put simply, measuring the behavioral consequences of an earlier act of inhibition on some target representation or process itself ought to engage the very same inhibitory processes to effect successful processing of that inhibited representation. If so, the putative cost of inhibition ought, to some degree, be mitigated by the putative benefit of inhibition at the time of test. When inhibitory functioning is impaired (either due to an inhibitory deficit, or to experimental manipulations) costs should be reduced, but so too should the benefits, opening the door for non-inhibitory components such as associative blocking to contribute to an effect. Thus, to effectively measure inhibition, a test should minimize the benefits of inhibition. We argued that the independent probe method, developed in the study of retrieval-induced forgetting (Anderson & Spellman, 1995) is such a test, and that the logic of this method can be applied in any domain in which inhibition is studied. This proposal has been validated in three series of studies using the think/no-think paradigm, which have consistently shown that the independent probe method provides a significantly more reliable index of inhibitory control deficits than do tests that suffer from the correlated costs and benefits problem. These findings indicate that the correlated costs and benefits problem seriously limits the ability to measure inhibition.

We would like to strongly emphasize that the correlated costs and benefits problem is not limited to memory retrieval, and applies broadly in domains such as executive control, visual attention, and language processing. We hope that recognition of this problem will help investigators interested in inhibition to develop theoretically targeted methods of testing inhibition theories, enabling the field to advance beyond an indefinite cycle of assertion and skepticism that plagues the study of such processes. We believe that sufficient evidence exists for inhibitory processes in memory retrieval, and that such processes are likely to contribute broadly to cognitive function, but reasonable doubt will remain as long as the field neglects the issue of measurement. Thus, our challenge to investigators in executive control, visual attention, and language, is to develop solutions to the correlated costs and benefits problem to see whether the properties of inhibition, identified in memory, also apply in other cognitive domains.

Figure Captions

Figure 1. A typical response-override situation. In this figure, a stimulus is associated with two responses, one of which is stronger (prepotent), and the other of which is weaker (dotted line). Response override occurs whenever one needs to either select the weaker, but more contextually appropriate response, or to simply stop the prepotent response from occurring. Inhibitory control is thought to achieve response override by suppressing activation of the prepotent response. This basic situation describes many paradigms in research on executive control, including the Stroop and go/no-go tasks.

Figure 2. A standard categorical retrieval-induced forgetting study. Illustrated here are two items from each of two categories that subjects have studied (typically 6 items are studied from 8 categories). In this example, subjects perform retrieval practice on "Fruit-Orange," but not on "Fruit-Banana" (unpracticed competitor) or on any members from the "Drinks" category (an unpracticed baseline category). The numbers show the percentage of items correctly recalled on the final cued-recall test. As shown here, practice facilitates recall of the practiced items relative to performance in baseline categories. Retrieval-induced forgetting is reflected in the reduced recall of unpracticed members of the practiced category (Banana), relative to performance in baseline categories (Scotch and Rum).

Figure 3. Final recall performance in the think/no-think procedure. The graph shows the percentage of items that subjects correctly recalled on the final test as a function of the whether they tried to recall the item (Respond), suppressed the item (Suppress), or had no reminders to the item during the think/no-think phase (Baseline). The left side shows recall when tested with the originally trained retrieval cue (i.e, the Same Probe), whereas

the right side shows recall when tested with a novel, extralist category cue (i.e., the Independent Probe). The numbers shown here were taken from a meta-analysis of 687 subjects run in the think/no-think paradigm in our lab.

Figure 4. fMRI results from Anderson et al (2004). Plotted above are the brain regions that significantly differed in activation between the Suppression trials and Respond trials during the think/ no-think phase (n=24). Areas in yellow were more active during Suppression trials than during Respond trials, whereas areas in blue were less active during Suppression ($P<0.001$). The white arrows highlight the reduced hippocampal activation in the Suppression condition.

Figure 5. Illustration of the correlated costs and benefits problem. A) In Situation A, a cue is associated with two responses. The inhibitory perspective predicts that practicing Item A in response to the cue, should result in inhibition of Item B (provided that it interferes). The inhibitory perspective also predicts in Situation B that practicing Item B should inhibit Item A, as Situation A and B are essentially equivalent. B) Illustrates how Situations A and B apply to the retrieval-practice paradigm. Practicing Orange during the retrieval practice phase should inhibit Banana. However, consistent application of the concept of inhibitory control at retrieval predicts that inhibition will also aid subjects in recalling Banana at the time of test by inhibiting Orange.

Figure 6. The correlated costs and benefits problem leads to unclear predictions on shared probe tests. Shown here is the amount of retrieval-induced forgetting on the final test that can be attributed to various causes as a function of inhibitory control ability. The solid black line shows the amount of RIF attributable to persisting suppression that was induced during retrieval practice, whereas the dotted line shows the amount attributable

to blocking by practiced competitors on the final test. Subjects with more inhibitory control should suffer more persisting suppression, whereas subjects with less inhibitory control should experience more blocking. The dashed line represents the total effect, considering the combined influences of these factors. This total changes little with inhibitory control ability, illustrating how same probe data may not clearly show differences between individuals with differing levels of inhibitory control ability.

Figure 7. Examples of potentially misleading conclusions based on shared probe data. Shown here are two studies that investigated RIF between different populations using only shared probe data. A) Moulin et al. (2002) found that both Alzheimer's patients and age-matched control subjects showed large RIF effects (recall of unpracticed items from practiced categories was lower than recall of Baseline items). B) Perfect & Fthenaki (2003) also found large RIF effects that did not differ across patients with frontal lobe damage, patients with temporal lobe damage, and control subjects. Both groups of investigators suggested that RIF may reflect a more automatic form of inhibition, distinct from inhibitory control. However, these groups would be expected to show same probe impairment if they were simply unable to prevent blocking from practiced competitors on the test.

Figure 8. The utility of the independent probe method, as illustrated with three different think/no-think studies. Each figure plots difference scores (baseline – suppression condition recall) to indicate the amount of memory impairment in the suppression condition. The columns show data from studies investigating the relation between inhibition and working memory (left column, Bell & Anderson, in preparation), age (middle column, Anderson, Reinholz, Mayr & Kuhl, in preparation), and trauma (right

column, Anderson & Kuhl, in preparation), respectively. In these studies, we predicted greater inhibition with greater working memory capacity, younger age, and higher level of experience with trauma. The data from all three studies showed the predicted pattern numerically in the same probe condition, but the differences in inhibition across groups were far from reliable. In contrast, the independent probe data clearly and significantly showed the predicted differences between the groups, and the three way interaction of group by inhibition by test type was significant in each series of studies.

References

Anderson, J. R. (1983). The architecture of cognition. Cambridge, MA: Harvard University Press.

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., & Bell, T.A. (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, 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). Academic Press.

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

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

Anderson, M.C. & Green, C. (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 theory. Journal of Experimental Psychology: Learning, Memory, & Cognition, 26, 1141-1159.

Anderson, M.C., & Kuhl, B.A. (in preparation). Psychological trauma and its enduring effects on memory suppression: Evidence for the plasticity of memory control.

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., Reinholz, J., Kuhl, B.A., & Mayr, U. (in preparation). Inhibition in aging and long-term memory: A cognitive aging study using the Think/No-Think paradigm.

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.

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. (2002). Semantic recall can cause episodic forgetting. Psychological Science, 13, 356–360.

Bell, T.A., & Anderson, M.C. (in preparation). Keeping things in, and out of mind: Individual differences in working memory capacity predict successful memory

suppression.

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: Lawrence Erlbaum Associates.

Bulevich, J.B., Roediger, H.L., Balota, D.A., & Butler, A.C. (in press). Failures to find suppression of episodic memories in the think/no-think paradigm. Memory & Cognition.

Camp, G., Pecher, D., & Schmidt, H.G. (2005). Retrieval-induced forgetting in implicit memory tests: The role of test awareness. Psychonomic Bulletin & Review, 12, 490-494.

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, M. A., & Fthenaki, A. (2003). Disruption of inhibitory control of memory following lesions to the frontal and temporal lobes. Cortex, 39, 667-686.

Depue, B.E., Banich, M.T., & Curran, T. (in press). Suppression of emotional and non-emotional content in memory: Effects of repetition on cognitive control. Psychological Science.

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.

Easton, P., & Gordon, P.E. (1984). Stabilization of Hebbian neural nets by inhibitory learning. Biological Cybernetics, 51, 1-9.

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., 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.

Gernsbacher, M.A. (1989). Mechanisms that improve referential access. Cognition, 32, 99-156.

Gernsbacher, M.A. and Faust, M.E. (1991) The mechanism of suppression: A component of general comprehension skill. Journal of Experimental Psychology: Learning, Memory & Cognition 17, 245-262.

Gernsbacher, M.A., Keysar, B., Robertson, R.R.W., & Werner, N.K. (2001). The role of suppression and enhancement in understanding metaphors. Journal of Memory and Language, 45, 433-450.

Gorfein, D. (Ed.). (2001). On the Consequences of Meaning Selection: Perspectives on Resolving Lexical Ambiguity. Washington, D.C.: American Psychological Association.

Hertel, P.T., & Calcaterra, G. (2005). Intentional forgetting benefits from thought substitution. Psychonomic Bulletin & Review, 12, 484-489.

Hertel, P.T., & Gerstle, M. (2003). Depressive deficits in forgetting. Psychological Science, 14, 573-578.

Hicks, J.L., & Starns, J.J. (2004). Retrieval-induced forgetting occurs in tests of item recognition. Psychonomic Bulletin & Review, 11, 125-130.

Hotta, C., & Kawaguchi, J. (2006, January). The effect of test delay on memory for suppressed items in the Think/No-Think paradigm. Poster presented at "The International Symposium on Inhibitory Processes in the Mind," Kyoto, Japan

Johnson, S.K., & Anderson, M.C. (2004). The role of inhibitory control in forgetting semantic knowledge. Psychological Science, 15, 448-453.

Joorman, J., Hertel, P.T., Brozovich, F., & Gotlib, I.H. (2005). Remembering the good, forgetting the bad: Intentional forgetting of emotional material in depression. Journal of Abnormal Psychology, 114, 640-648.

Kawaguichi, J., & Hotta, C. (2006,January). Intention to forget, but not distraction is critical to forget: the effect of distraction task on memory retrieval in a five minute delay test. Poster presented at "The International Symposium on Inhibitory Processes in the Mind," Kyoto, Japan.

Levy, B.J., & Anderson, M.C. (2002). Inhibitory processes and the control of memory retrieval. Trends in Cognitive Sciences, 6, 299-305.

Levy, B.J., McVeigh, N.D., Marful, A., & Anderson, M.C. (under review). Inhibiting your native language: The role of retrieval-induced forgetting during second language acquisition.

MacLeod, M.D. & Saunders, J. (2005). The role of inhibitory control in the production of misinformation effects. Journal of Experimental Psychology: Learning, Memory & Cognition, 31, 964-979.

Mayr, U., & Keele, S.W. (2000). Changing internal constraints on action: The role of backward inhibition. Journal of Experimental Psychology: General, 129, 4-26.

Mayr, U. (2001). Age differences in the selection of mental sets: The role of inhibition, stimulus ambiguity, and response-set overlap. Psychology and Aging, 16, 96-109.

Mayr, U., Diedrichsen, J., Ivry, R., & Keele, S.W. (in press). Dissociating task-set selection from task-set inhibition in prefrontal cortex. Journal of Cognitive Neuroscience.

Menon, V., Adleman, 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.

Milliken, B., Joordens, S., Merikle, P.M., & Seiffert, A.E. (1998). Selective attention: A reevaluation of the implications of negative priming. Psychological Review, 105, 203-229.

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.

Neill, W.T., & Valdes, L.A. (1996). Facilitatory and inhibitory aspects of attention. In A.F. Kramer, M.G.H. Coles, & G.D. Logan (Eds.), Converging operations in the study of visual selective attention (pp. 45-76). Washington, DC: American Psychological Association.

Phenix T.L., & Campbell J.I.D. (2004). Effects of multiplication practice on product verification: Integrated structures model or retrieval-induced forgetting? Memory & Cognition, 32, 324-335.

Radvansky, G.A. (1999). Memory retrieval and suppression: The inhibition of situation models. Journal of Experimental Psychology: General, 128, 563-579.

Raaijmakers, J.W., & Shiffrin, R.M. (1981). Search of associative memory. Psychological Review, 88, 93-134.

Shivde, G., & Anderson, M.C. (2001). The role of inhibition in meaning selection: Insights from retrieval-induced forgetting. In D. Gorfein (Ed.), On the Consequences of Meaning Selection: Perspectives on Resolving Lexical Ambiguity (pp. 175-190). Washington, D.C.: American Psychological Association.

Simpson, G.B. and Kang, H. (1994) Inhibitory processes in the recognition of homograph meanings. In Inhibitory Processes in Attention, Memory, and Language (Dagenbach, D. and Carr, T.H., eds), pp. 359-381, Academic Press.

Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. Psychological Review, 99, 195-231.

Starns, J.J., & Hicks, J.L. (2004). Episodic generation can cause semantic forgetting: Retrieval-induced forgetting of false memories. Memory & Cognition, 32, 602-609.

Tipper, S.P. (2001). Does negative priming reflect inhibitory mechanisms? A review and integration of conflicting views. The Quarterly Journal of Experimental Psychology, 54A, 321-343.

Veling, H., & Van Knippenberg, A. (2004). Remembering can cause inhibition. Retrieval-induced inhibition as a cue independent process. Journal of Experimental Psychology: Learning, memory, & cognition. 30, 315-318.

Wessel, I., Wetzels, S., Jelicic, M., Merckelbach, H. (2005). Dissociation and memory suppression: A comparison of high and low dissociative individuals' performance on the Think-No Think task. Personality and Individual Differences, 39, 1461-1470.

Zellner, M., & Bäuml, K. (2005). Intact retrieval inhibition in children's episodic recall. Memory & Cognition, 33, 396-404.

Figure 1

Typical Response Override Situation

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

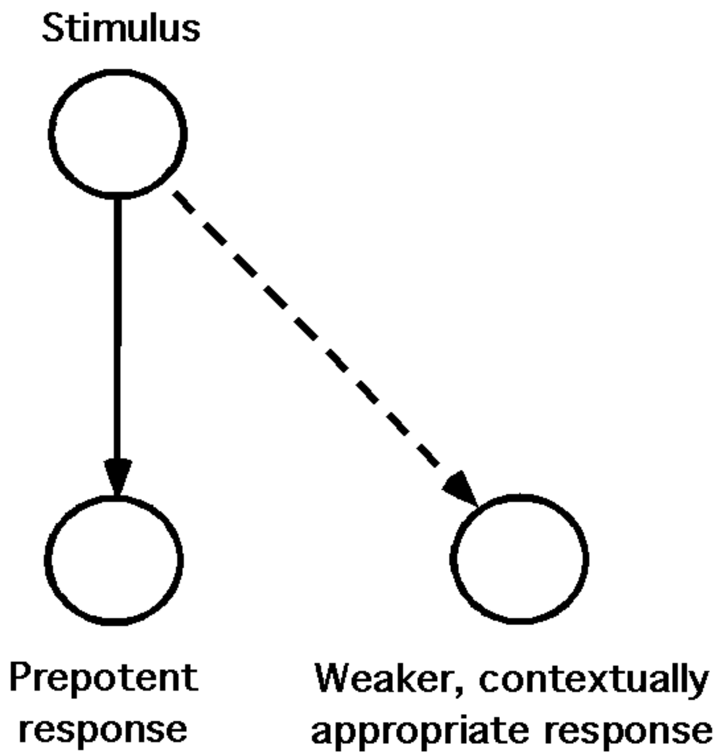

Figure 2

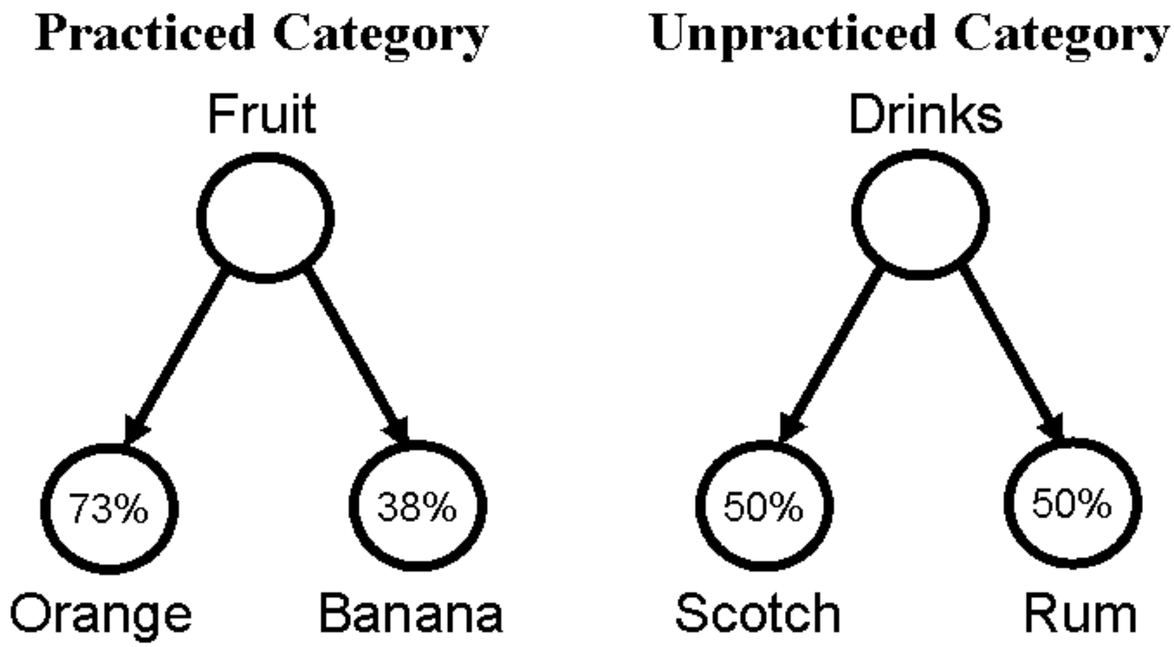
Practiced Category
Fruit
73%
38%
Orange
Banana
Unpracticed Category
Drinks
50%
50%
Scotch
Rum

Figure 3

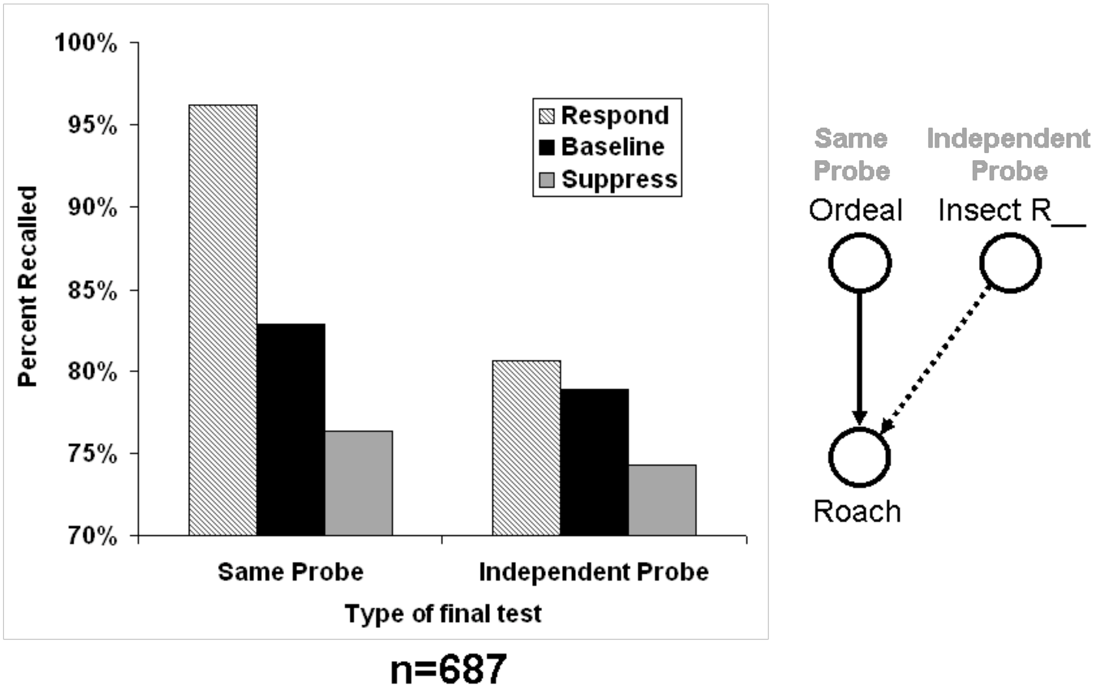
100%
95%
90%
85%
80%
75%
70%
Percent Recalled
Respond
Baseline
Suppress
Same Probe
Independent Probe
Type of final test
n=687
Same Probe
Independent Probe
Ordeal
Insect R__
Roach

Figure 4

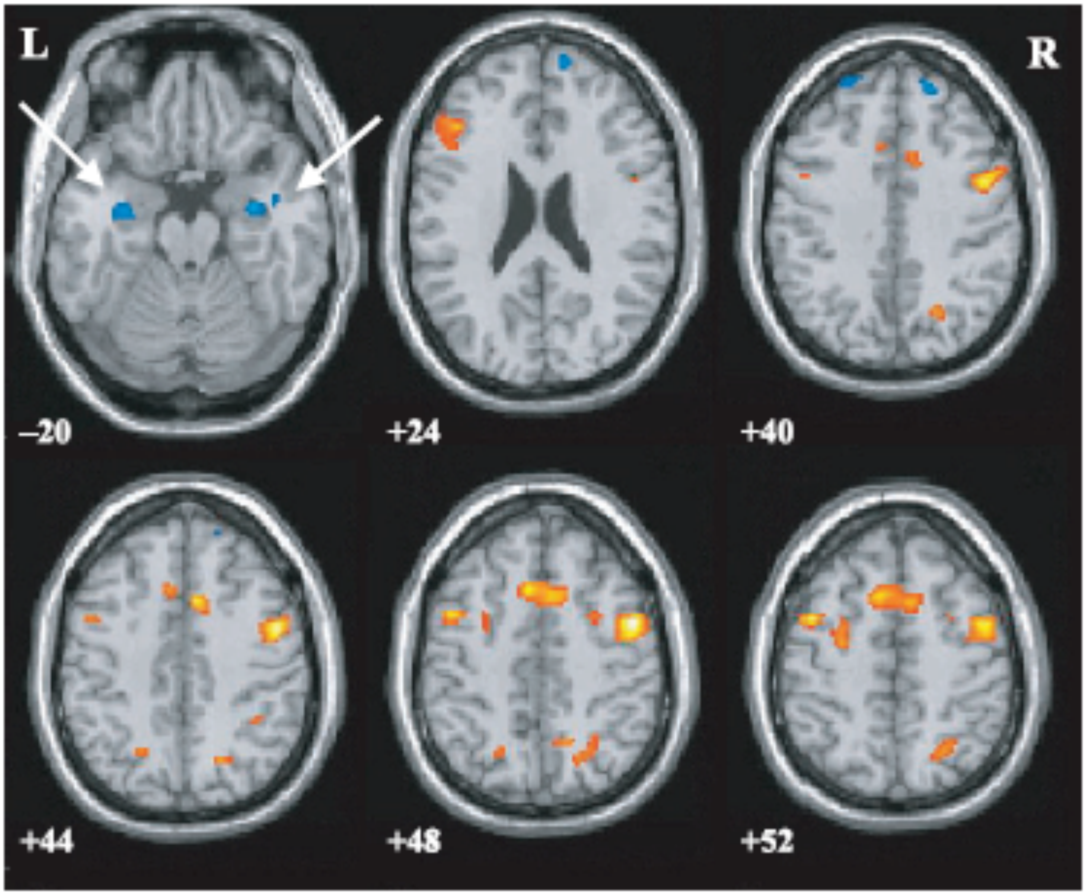

Figure 5

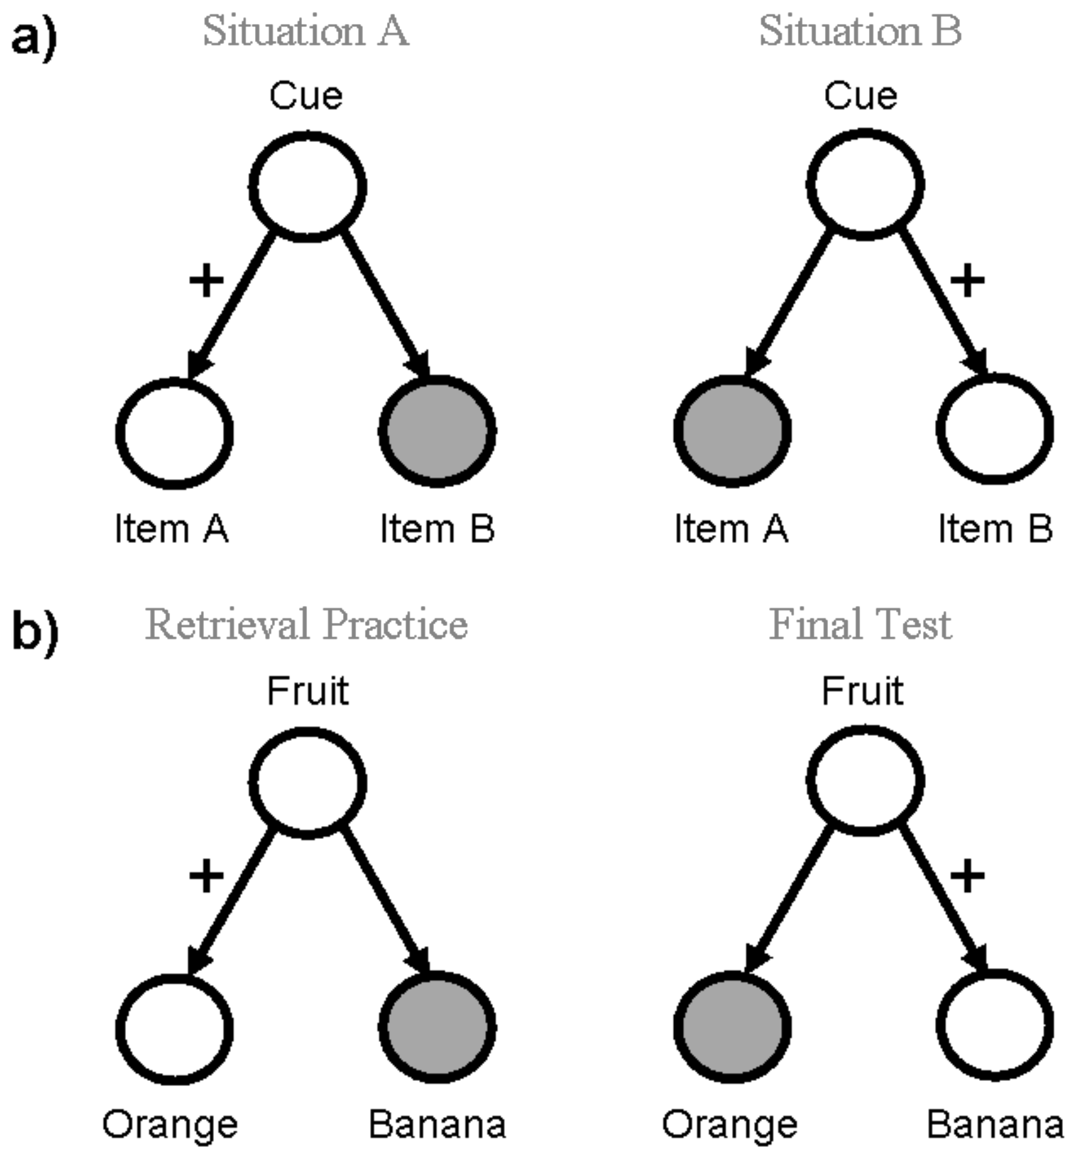
a)
Situation A
Cue
+
Item A
Item B
Situation B
Cue
+
Item A
Item B
b)
Retrieval Practice
Fruit
+
Orange
Banana
Final Test
Fruit
+
Orange
Banana

Figure 6

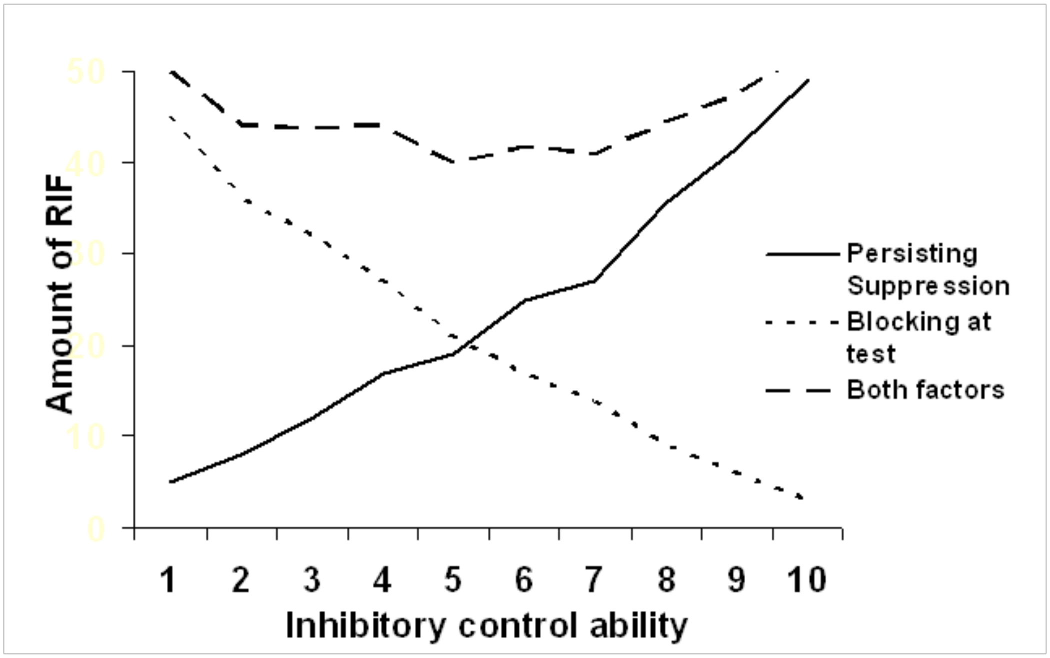
Amount of RIF
1
2
3
4
5
6
7
8
9
10
Inhibitory control ability
Persisting Suppression
Blocking at test
Both factors

Figure 7

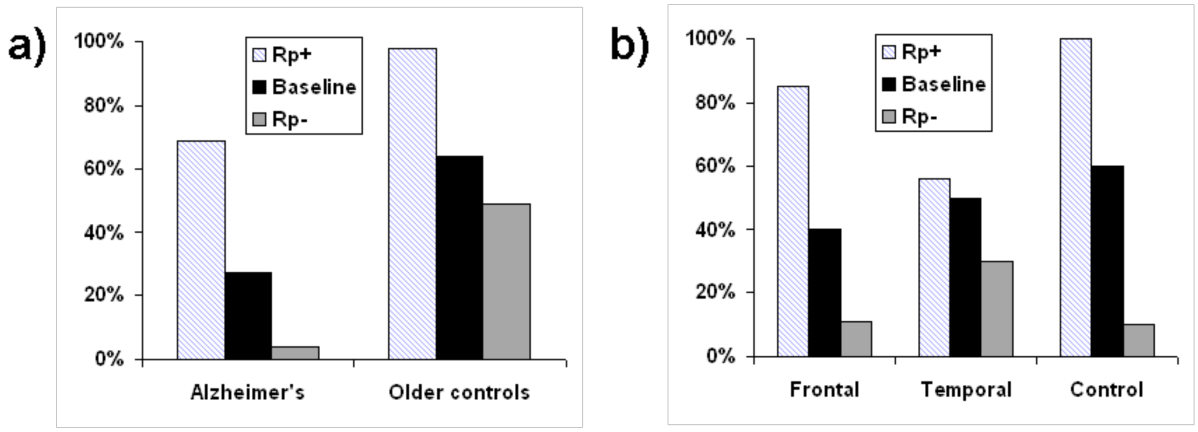
a)
100%
80%
60%
40%
20%
0%
Rp+
Baseline
Rp-
Alzheimer's
Older controls
b)
100%
80%
60%
40%
20%
0%
Rp+
Baseline
Rp-
Frontal
Temporal
Control

Figure 8

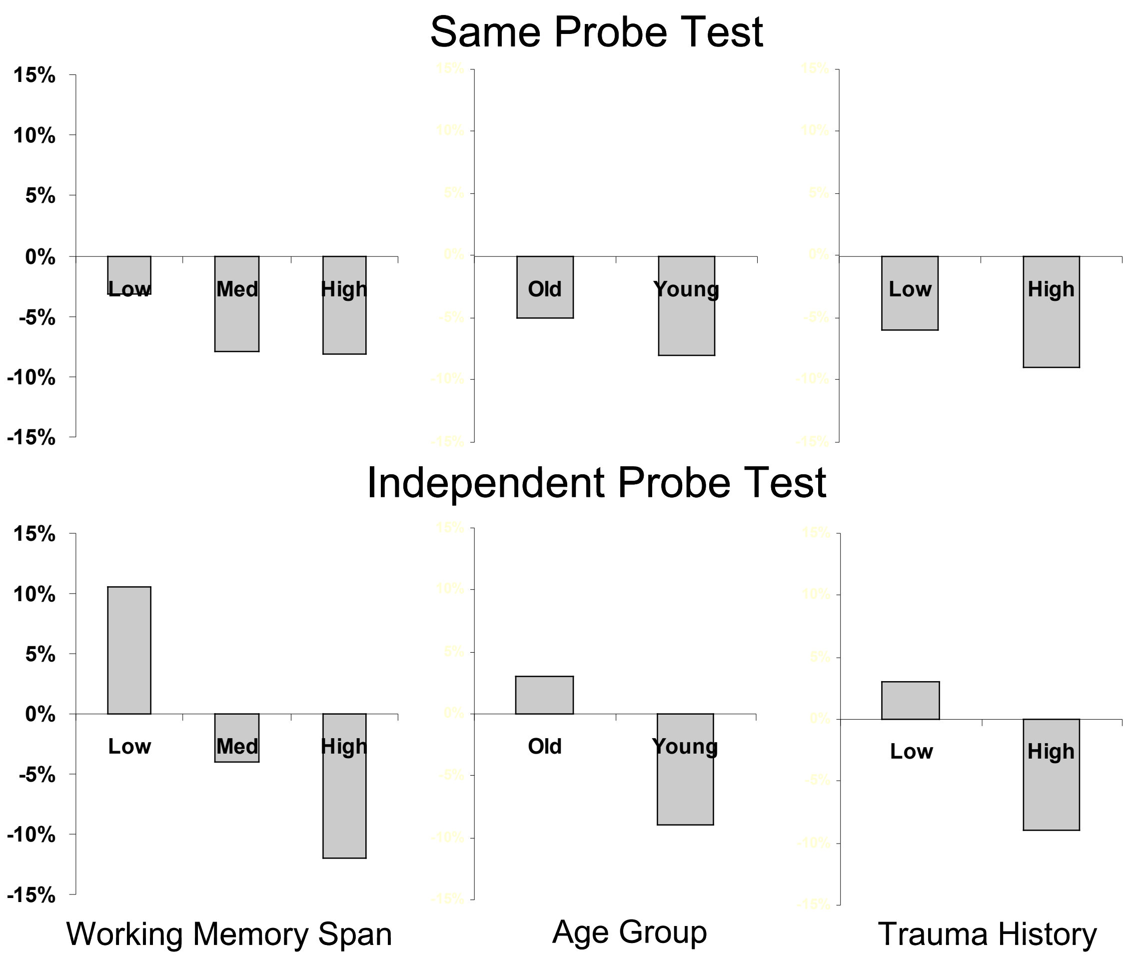
Same Probe Test
15%
10%
5%
0%
-5%
-10%
-15%
Low
Med
High
Old
Young
Low
High
Independent Probe Test
15%
10%
5%
0%
-5%
-10%
-15%
Low
Med
High
Old
Young
Low
High
Working Memory Span
Age Group
Trauma History