# 6

# ON THE RELATIONSHIP BETWEEN INTERFERENCE AND INHIBITION IN COGNITION

Michael C. Anderson and Benjamin J. Levy

There is a moment that each of us knows all too well. This moment arises when, by chance, we encounter something that reminds us of an experience we would rather not think about. For some, this might occur when opening a drawer to see an object given to us by a loved one that has been lost due to death or a broken relationship. For others—such as for veterans returning from combat—existence may be an ongoing battle against their own memories, and against intrusive recollections of wartime experiences. Whatever the trigger, this moment elicits certain key experiences: an abrupt feeling of alarm and arousal, followed quickly by an effort to shut down retrieval, regain footing in our mental landscape, and redirect attention to more profitable goals. Even neutral memories can be too accessible; anyone who has ever gone to yesterday's parking spot, used an outdated bank code, or called out to one's previous spouse instead of one's current one knows well the trouble such interloping knowledge can induce. Thus, one of the most innocent assumptions one might make about memory—that remembering is good and forgetting is bad—is, in fact, often untrue. More often than we realize, forgetting is our goal, and remembering is the human frailty.

Given that forgetting can at times be a positive goal, it becomes important to address how we accomplish it, when successful. Although there are likely to be a variety of ways people limit the accessibility of unwanted memories, there has been considerable emphasis on the role of inhibition in accomplishing this function. Inhibition refers to a hypothesized control mechanism that reduces the level of activation associated with a trace, rendering it less accessible to ongoing cognition. The existence of inhibitory processes has been hypothesized in many different domains of human cognition, including attention, memory, language, and motor action (see Dagenbach & Carr, 1994; Dempster & Brainerd, 1995; Gorfein & MacLeod, 2007). In memory research, the inhibition view hypothesizes that when the activation associated with a trace is disruptive, inhibition may be engaged to reduce this unwanted accessibility. The persistence of this inhibition is thought to induce forgetting. Thus, inhibition may be one instrumental process in achieving adaptive or *functional forgetting.*

In this chapter, we consider inhibition as a response to unwanted accessibility. First, we review the evidence for a relationship between the degree of interference caused by a trace and inhibitory aftereffects observed on later retention tests for that trace. This evidence is highly consistent with the idea that inhibitory mechanisms are triggered in response to intrusive memories that might otherwise disrupt our current goals. As such, inhibition appears to be a key mechanism for achieving functional forgetting. We then consider the theoretical relationship between interference and inhibition, and what this relationship predicts about the expected aftereffects of inhibition. We introduce the concept of a demand-success trade-off, a crucial factor complicating the measurement of inhibition. Demand-success trade-offs have important consequences for testing theoretical models of inhibition as well as theories that posit inhibitory deficits in different populations. These trade-offs are not unique to memory research, but affect any domain concerned with inhibition. By drawing attention to this issue, we hope to prevent confusion and theoretical controversy in the literature on inhibition, and improve the assessment of inhibitory deficits in special populations.

## EVIDENCE FOR A ROLE OF INHIBITION IN COUNTERING UNWANTED ACCESSIBILITY

Though we may often complain about our memories, they can be surprisingly efficient. A stimulus can remind us of a related experience with little effort or intention. As useful as this ability is, it is less helpful when our goal is to think of something other than the first memory that leaps

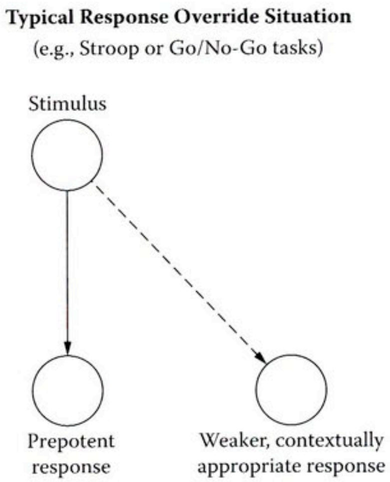


**Figure 6.1** A typical response override situation. In this figure, a stimulus is associated with two responses, one of which is stronger (prepotent), and the other of which is weaker (dotted line). Response override occurs whenever one needs either to select the weaker, but more contextually appropriate, response, or simply to 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.

to mind. Sometimes, unwanted retrievals arise while we are retrieving something else, impeding the recall of the desired memory. Other times, we don't mean to recall anything at all; automatic retrieval in response to an unexpected reminder may elicit unpleasant or otherwise distracting memories. When the most accessible trace is unwelcome, we must engage a control process to overcome interference posed by these overly zealous memories simply to reorient to the desired task (see Figure 6.1). A core claim of the adaptive forgetting view is that these situations—when accessibility is disruptive—trigger inhibitory processes that suppress the intruding representation, diminishing its accessibility, helping us to regain control over our thoughts. In this section, we describe studies examining this key assumption of the adaptive forgetting hypothesis in two situations: selective retrieval and retrieval stopping.

### *Interference-Dependence During Selective Retrieval*

In selective retrieval situations, the goal is to recall a particular memory given a cue that activates other competing traces. To the extent that nontarget memories interfere with retrieval of the target, inhibitory processes ought to be engaged. The role of inhibitory control during selective retrieval is often studied using the retrieval practice procedure (Anderson, Bjork, & Bjork, 1994). In a typical version of this paradigm participants study multiple exemplars of several categories (e.g., *fruits-banana, fruits-lemon, tools-screwdriver, tools-hammer*). A retrieval practice phase follows in which participants are cued to recall some of the exemplars from some of the categories (e.g., *fruits-ba*______). Afterwards, they receive a final test in which they are cued to recall all studied items. Importantly, some of the originally studied categories do not appear during the retrieval practice phase, and thus serve as a baseline for how memorable items should be on this final test, given that they have received no intervening practice. As one might expect, the practiced items are recalled more often than baseline items, confirming that retrieval practice is effective at boosting memory (Bjork, 1975). But what happens to unpracticed exemplars in practiced categories (e.g., *fruits-lemon*)? Our earlier analysis suggests that the category cues during retrieval practice may remind subjects not only of the desired item, but also of the other exemplars. Thus, these unpracticed items may interfere during retrieval practice, triggering inhibition as a means of resolving competition. If so, these items may be harder to recall on the final test than the baseline items. In fact, this pattern is typically observed in the retrieval practice paradigm (see Figure 6.2). This memory deficit, known as *retrieval-induced forgetting* (RIF), is consistent with an inhibitory aftereffect induced by selective retrieval that undermines retention of competing items (for a more extensive review of RIF, see Chapter 5, this volume).

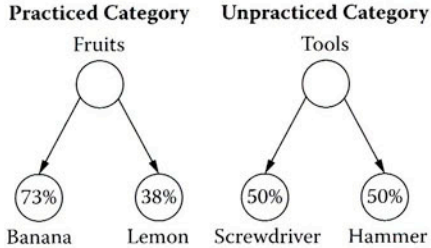


**Figure 6.2** A standard categorical RIF study. Illustrated here are two items from each of two categories that subjects have studied (typically six items are studied from eight categories). In this example, subjects perform retrieval practice on *fruits-banana*, but not on *fruits-lemon* (unpracticed competitor) or on any members from the "tools" 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. RIF is reflected in the reduced recall of unpracticed members of the practiced category (*banana*), relative to performance in baseline categories (*screwdriver* and *hammer*).

If RIF truly reflects an inhibitory aftereffect that is triggered by competition, we should find that the magnitude of this effect is larger when interference needs to be resolved than when related items are not interfering. The dependency of RIF on the need to resolve interference, known as *interference-dependence*, has been examined in several ways, using manipulations of stimulus materials and retrieval task.

*Manipulations of Competitor Interference* One approach to examining interference-dependence involves manipulating the prepotency of competing items by selecting materials that vary in the strength of the cue-to-target association. For example, *banana* and *orange* are higher-frequency examples of the category *fruits* than are *kiwi* or *papaya*. Thus, high-frequency items should be easier to recall given *fruits* as a cue and, for the same reasons, should interfere more when retrieving other fruits. If so, high-frequency exemplars should be more likely to trigger inhibition, and produce larger RIF effects. In contrast, low-frequency exemplars should be harder to retrieve given their category. If so, they should cause less interference and suffer little RIF. As expected, high-frequency items are recalled better overall than low-frequency items, and critically, RIF is reliably greater for those high-frequency items (Anderson et al., 1994), consistent with interference-dependence.

Manipulations of the prepotency of a competitor are not limited to taxonomic categories. A similar manipulation can be achieved using homographs in which one meaning of the word is far more dominant than the other. With materials such as these, one's strong tendency, upon reading the word, is to retrieve the dominant meaning, even if one knows that it is inappropriate. For example, when asked to retrieve an associate of the verb meaning of the following words, it is difficult to avoid thinking of their noun meanings first: *root*, *prune*, *fence*, *lobby*, and *stump*. Because the accessibility of the noun meaning interferes with the intended task, one might expect inhibitory processes to be engaged to inhibit the dominant sense and facilitate retrieval of the subordinate meaning. This is indeed the pattern that is observed (Johnson & Anderson, 2004; Shivde & Anderson, 2001). In contrast, retrieving the dominant meaning is effortless because the subordinate meaning poses no interference, and as expected, this produces no inhibition of the weaker meaning (Shivde & Anderson, 2001).

Another common instance of unwanted accessibility occurs during foreign language learning. Learning a new language requires that we attach entirely new phonological labels to objects and ideas that already have strong associations to native language words. As such, when we try to produce the foreign name for an object, we must overcome unwanted activation from the native language name. Thus, we might expect that retrieving words in a foreign language should trigger inhibition, especially during the early stages of learning. Indeed, for beginning Spanish speakers, when new words are not yet well learned, this is exactly the pattern observed: Subjects have difficulty accessing specific verbal labels from their native language after repeatedly retrieving the corresponding label in the second language (Levy, McVeigh, Marful, & Anderson, 2007). This pattern depends on participants' fluency in their second language: Participants with strongly asymmetric skills (e.g., those who know English much better than Spanish) show sizable inhibition effects, whereas those with more symmetric skills (i.e., show evidence of comparable fluency) show little inhibition. Thus, when retrieval of foreign language words was difficult, inhibition suppressed native language words, consistent with an adaptive forgetting process. This dynamic may form the basis of the puzzling phenomenon of first language attrition, or the forgetting of one's native language when immersed in a foreign language (e.g., Linck, Kroll, & Sunderman, 2009).

Most manipulations of competitor interference have varied the strength of competition by manipulating whether competitors have a low or high a priori association to the retrieval practice cue. One notable exception to this trend is a study conducted by Storm, Bjork, and Bjork (2007), who manipulated the interference properties of competitors by a procedural manipulation. On each trial, before retrieval practice occurred, subjects were presented with short lists of category-exemplar pairs to study (e.g., *country-Russia*, *flower-lily*, *country-Sweden*, *flower-tulip*). A directed forgetting manipulation then instructed subjects to either forget or remember these items. Prior work suggests that this instructional manipulation is effective in varying how accessible the prior items are (e.g., Bjork, 1970), and this assumption was independently verified. This directed forgetting cue was followed by semantic generation (i.e., retrieval practice) of other items from one of those categories (e.g., *flower-pa____* for *pansy*). Thus, the items from the list studied immediately prior to retrieval practice acted as the competitors during semantic generation, and the instructional cue (insofar as it was effective) manipulated the degree to which these items would interfere. Strikingly, cuing subjects to forget eliminated RIF for those competing items, whereas RIF was observed for items they were instructed to remember. Storm et al. argued that directed forgetting instructions made these exemplars less accessible, providing less interference during retrieval practice, and therefore they did not trigger inhibitory processes. Thus, even when participants studied the same item and performed the same retrieval practice task, RIF could be made to appear or disappear, depending on how accessible those items were entering the retrieval practice phase.

*Manipulations of Practice Task Interference Demands* Another approach to assessing interference-dependence is to manipulate the demands of the retrieval practice task, instead of the strength of the competitor. If a cue activates a nontarget memory, but that activation is unlikely to undermine success on the retrieval practice task, then inhibitory processes should be unnecessary.

The most straightforward version of this approach comes from studies that contrast the effects of retrieval practice with those of repeated reexposure to the same stimuli (often referred to as extra presentations). Here all aspects of the retrieval practice paradigm are matched across two groups of subjects, except for the events during the practice phase. One group performs retrieval practice trials, as usual (e.g., *fruit-or_____* for *orange*), while the reexposure group is given the entire category-exemplar pair for additional study (e.g., *fruit-orange*). In this design, both groups practice the same exemplars an equal number of repetitions for the same amount of time, but the extra exposures group has the target item fully specified. Importantly, the inhibitory account predicts that, to the extent that reexposure poses very little demand on interference resolution, it should not produce forgetting on the final test. In contrast, noninhibitory explanations such as associative blocking predict that forgetting should occur regardless of how the practiced items were strengthened (see Anderson & Bjork, 1994, for further discussion of noninhibitory accounts of RIF). Many studies using this approach have found RIF after retrieval practice, but not after reexposure, provided that output interference is controlled at test (Bäuml, 1996, 1997, 2002; Johansson, Aslan, Bäuml, Gabel, & Mecklinger, 2007; Saunders, Fernandes, & Kosnes, 2009; Wimber, Rutschmann, Greenlee, & Bäuml, 2009). The dependency of RIF on active retrieval generalizes to retrieval of visuospatial information (Ciranni & Shimamura, 1999), homograph meanings (Shivde & Anderson, 2001), propositions (Anderson & Bell, 2001), and arithmetic facts (Campbell & Phenix, 2009), suggesting it is a general attribute of this phenomenon.

Even when two groups receive retrieval practice, however, the difficulty of the interference resolution task can be manipulated by changing the retrieval task itself. For instance, Anderson, Bjork, and Bjork (2000) required all subjects to engage in retrieval practice, but simply manipulated whether the practice required the resolution of interference. The competitive group was provided with the category name and a letter stem for the exemplar (e.g., *fruit- or____*) as is customarily done; the noncompetitive group was instead asked to recall the category name, given the exemplar (e.g., *fr___-orange*). Although both tasks engage retrieval, the latter specifies the category as the target, so other exemplars should pose no threat to retrieval success and no RIF should be observed. Consistent with the adaptive inhibition view, competitive, but not noncompetitive, retrieval practice induced RIF, despite comparable strengthening of practiced items between the two groups.

Another clever approach for manipulating interference demands was introduced by Bajo, Gomez-Ariza, Fernandez, and Marful (2006). During retrieval practice, they presented the shared category cue either two seconds before giving the subjects the individuating letter stem for the exemplar or presented them simultaneously, as is conventionally done. By presenting the category two seconds before the individuating stem, Bajo et al. hoped to increase the likelihood that competitors would be activated enough to pose interference when the distinctive stem was provided. If so, greater RIF should be found than in a case where the distinctive stem is provided immediately with the category. Consistent with this, they observed significantly more RIF when the category preceded the distinctive stem, despite comparable retrieval practice success and comparable strengthening of practiced items in both cases. Interestingly, these findings were observed with a paradigm in which categories were defined purely on the basis of lexical, rather than semantic, attributes (each category consisted of words that shared the same initial two-letter syllable and differed in their third letter). Thus, they were able to manipulate the amount of interference provided by these lexical competitors through the use of two cuing methods that differed only in their timing and not in objective information provided.

If inhibition is recruited to facilitate retrieval, what happens if the retrieval task itself is impossible to accomplish? Is RIF still found for competing items, even if nothing is retrieved during retrieval practice? To look at this, Storm and colleagues (e.g., Storm, Bjork, Bjork, & Nestojko, 2006; Storm & Nestojko, 2010) designed their stimuli so that retrieval practice success was impossible for some trials by providing letter stems that did not correspond to any appropriate exemplar. Even in these situations where retrieval failure is ensured for every retrieval practice trial from a specific category, subjects show robust RIF. These findings suggest that successful retrieval itself is not a prerequisite for observing RIF and that inhibition functions to assist difficult retrieval (see also Dagenbach, Carr, & Barnhard, 1990). Intriguingly, on the retrieval practice trials that were possible to complete, Storm et al. (2006) noted that the subjects with lower practice success rates were the ones who suffered the most RIF, indicating that when retrieval practice was difficult, inhibition was more pronounced.

Functional neuroimaging studies provide converging evidence that inhibition reflects a cognitive control response to unwanted accessibility (for a review, see Levy, Kuhl, & Wagner, in press). It is known that the prefrontal cortex is involved in cognitive control generally (e.g., Miller & Cohen, 2001), and specifically in response to interference in memory (Dolan & Fletcher, 1997; Fletcher, Shallice, & Dolan, 2000; Henson, Shallice, Josephs, & Dolan, 2002; Shimamura, Jurica, Mangels, Gershberg, & Knight, 1995). With this in mind, Wimber et al. (2009) used functional magnetic resonance imaging (fMRI) to contrast activity during blocks of competitive retrieval practice with activity during extra presentations of the same items, as had been done in the numerous behavioral studies described above. Strikingly, competitive retrieval was associated with increased activity within the prefrontal cortex, and the extent of this increased activity predicted later RIF for competing items. Kuhl, Dudukovic, Kahn, and Wagner (2007) also found that retrieval practice engaged the prefrontal cortex (PFC) and, interestingly, prefrontal regions showed a pattern of decreasing activation across repeated retrieval practice trials. Thus, control-related regions are strongly engaged initially when unwanted accessibility is greatest, and become less involved as control demands diminish with practice. Indeed, subjects who showed the largest decrease in prefrontal activity across retrieval practice attempts were the ones who showed the most RIF. These correlations between PFC and RIF provide converging evidence for the view that interference experienced during selective retrieval triggers inhibition as an adaptive forgetting process.

### *Interference-Dependence During Retrieval Stopping*

Although we have focused thus far on selective retrieval, people encounter intrusive memories in many situations. Sometimes, for example, cues bring to mind memories even when we are not seeking to retrieve something, and these memories may be disruptive. When uninvited memories intrude, people often attempt to suppress the retrieval process and exclude the intrusive memory from awareness. The effort to regulate unwanted accessibility in these situations ought to recruit inhibitory mechanisms that facilitate adaptive forgetting of the disruptive event.

Intentional retrieval suppression has been studied with the think/no-think paradigm (Anderson & Green, 2001). In these studies subjects learn pairs of items in which one item acts as a cue to remind them of a particular response (e.g., *lawn-beef*, *journey-pants*). Later they are shown these cues again in isolation (e.g., *lawn*) and are asked to engage control over retrieval. On think trials subjects are asked to remember the response (*beef*) and keep it in mind during the trial. On no-think trials, they are asked to attend to the cue (*journey*), but to prevent the associated word (*pants*) from entering awareness. After performing these tasks on different sets of cues, subjects are tested on all of the previously studied responses. Unsurprisingly, response words that were practiced during the think trials are recalled more often than baseline items that were initially learned, but that did not occur during the intervening think/no-think phase (see Figure 6.3). The no-think response words, however, are recalled less often than the baseline items. The impaired recall of no-think items suggests that they have been disrupted by control mechanisms (Anderson & Green, 2001; Anderson et al., 2004; Depue, Banich, & Curran, 2006; Depue, Curran, & Banich, 2007). Importantly, no-think responses (e.g., *pants*) are harder to recall even when subjects are tested with novel extralist cues, referred to as independent probes (Anderson & Green, 2001; Anderson et al., 2004), suggesting that the memory deficit arises from reductions in the accessibility of the response (Anderson & Spellman, 1995).

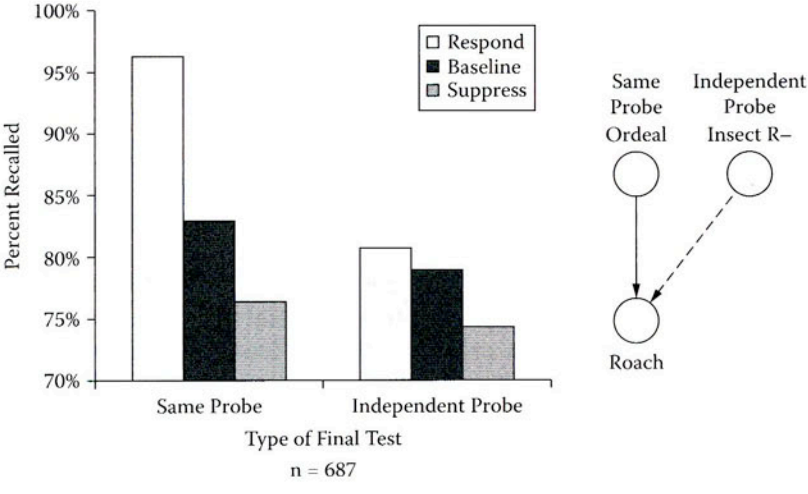


**Figure 6.3** Final recall performance in the think/no-think procedure. The graph shows the percentage of items that subjects correctly recalled on the final test as a function of whether they tried to recall the item (respond), suppressed the item (suppress), or had no reminders to the item during 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.

In recent work using this paradigm, we sought to study the relationship between the amount of memory inhibition people show on a final retention test and how effectively they reduced the intrusiveness of the unwanted memory during retrieval suppression itself. According to the adaptive forgetting hypothesis, evidence for inhibitory forgetting should be most evident for people who are effective at limiting the intrusiveness of an unwanted memory during suppression. To study this relationship, we developed an online, trial-by-trial measure of how intrusive an item is during no-think trials (Levy & Anderson, in preparation). After each no-think trial subjects reported whether or not the response word came to mind, allowing us to identify when conscious intrusions happened and to track their progression with practice at the suppression task. Thus, in contrast to studies of RIF in which intrusiveness is presumed on the basis of a competitor's a priori associative strength to a cue, this study directly assessed people's subjective experience of whether the unwanted memory intruded.

Using this metric of intrusiveness, we were able to document several important findings about intrusions and memory inhibition. First, when people attempt to suppress retrieval of an unwanted memory, they often fail. Indeed, on the initial suppression attempt for a given item, people usually fail to prevent the item from coming to mind. With repeated practice, however, intrusions decrease in frequency, showing that people can, with repeated effort, limit accessibility of unwanted memories (see Figure 6.4). Intriguingly, this steady decline in intrusions resembles

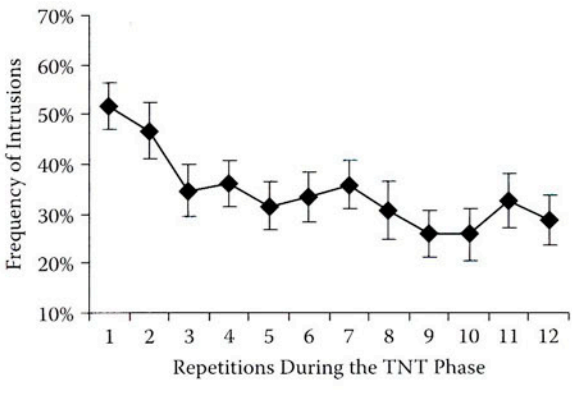


**Figure 6.4** Frequency of intrusions during a think/no-think study. Plotted are the percentage of no-think trials where the subject reported that the response word came to mind for each of the 12 repetitions of the no-think items. Subjects initially report intrusions on more than half of the trials, but the frequency of these intrusions rapidly declines with practice. (From Levy & Anderson, in preparation.)

the memory impairment observed for no-think items, which gradually increases with repeated practice at the task (Anderson & Green, 2001). This similarity suggests that subjective intrusion ratings and final test performance are measuring the same underlying construct: the state of activation of the to-be-avoided memory. If so, subjects who show the most forgetting of suppressed items on the final test (i.e., the good inhibitors) should also show the steepest decline in intrusion frequency during the think/no-think task itself. Consistent with this, good inhibitors showed a much more robust decline in intrusions over repetitions than poor inhibitors, as predicted by adaptive forgetting.

The measure of memory intrusiveness described above provides a valuable tool in identifying the neural mechanisms of adaptive forgetting, and the conditions under which they operate. For instance, incorporating the intrusion measure into a functional magnetic resonance imaging design has allowed us to examine how the brain responds to an intrusion. Prior functional neuroimaging data have suggested that mnemonic control is achieved in the think/no-think task by downregulating activity within the hippocampus (Anderson & Levy, 2009; Anderson et al., 2004; Depue et al., 2007; ), a region known to be active during memory retrieval. Using the intrusion paradigm described above, Levy and Anderson (in preparation) found that the downregulation of hippocampal activity is significantly greater during suppression trials on which subjects briefly experience an intrusion than when they do not. Indeed, the magnitude of hippocampal downregulation in reaction to intrusions predicted subsequent forgetting for suppression items on the final test. This finding provides intriguing support for the idea that controlled modulation of memory structures in the brain underlies the ability to suppress retrieval, and that this response reacts directly to the experience of unwanted accessibility. Thus, in retrieval stopping, as in retrieval selection, unwanted accessibility triggers a modulatory response that reduces accessibility, consistent with the adaptive forgetting hypothesis.

## THE DEMAND-SUCCESS TRADE-OFF PROBLEM: A CENTRAL THEORETICAL ISSUE FOR RESEARCH ON INHIBITORY CONTROL

As the foregoing findings illustrate, excess activation of an interfering trace triggers processes that impair retention of that trace. Findings consistent with this conclusion arise regardless of whether one manipulates the characteristics of the interfering trace or the retrieval task itself.

Collectively, these findings support the view that inhibitory processes are engaged to suppress excess activation, consistent with an adaptive forgetting mechanism.

Given the evidence for interference-dependence, one might surmise that behavioral indices of inhibition should grow as a function of how intrusive a memory is. After all, inhibition functions to overcome interference, so this prediction might seem obvious, especially given the prior findings. Indeed, inhibition theories in every domain—language, attention, memory, or motor action—hypothesize that inhibition emerges as a response to excess activation, and so behavioral measures of the aftereffects of inhibition should increase with interference.

Nevertheless, the predicted relationship between interference and measured inhibitory aftereffects is not straightforward. Depending on how interference is manipulated, behavioral indices of inhibition may increase or decrease with increasing interference. This complexity stems from a theoretically predicted relationship we refer to as a *demand-success trade-off*, which arises whenever an inhibitory theory assumes that inhibition is anything less than perfectly effective. Here, we describe the essence of the demand-success trade-off problem, providing several illustrations. We argue that this problem poses a fundamental challenge for using behavioral aftereffects of inhibition to infer properties of the mechanism or to measure variations in that mechanism across populations. Importantly, this problem is not limited to memory, but affects all domains concerned with inhibition.

### *The Essence of the Demand-Success Trade-Off Problem*

Since the putative function of inhibition is to reduce the influence of an unwanted representation, the need for inhibition should be related to how intrusive an unwanted trace is. Indeed, nearly every inhibition theorist would agree that as competitors become more interfering, the *demand* for inhibition should rise, increasing the likelihood that inhibition is triggered. Thus, inhibition theories might seem to predict a positive relationship between interference and behavioral aftereffects of inhibition.

A problem arises, though, because inhibition may be imperfect. Indeed, much of the interest in inhibitory control emanates from the potential of this construct to explain individual differences in handling interference. Prominent theories of cognitive aging, cognitive development, and clinical disorders, including depression, schizophrenia, obsessive compulsive disorder, and attention-deficit disorder, hypothesize deficits in inhibitory control functions, subserved by the prefrontal cortex, that underlie a broad spectrum of cognitive deficits that influence attention, language, working memory, and long-term memory. Indeed, inhibitory control has even been suggested as an important process in general intelligence. This emphasis on variable inhibitory abilities is reasonable, given increased sensitivities to interference exhibited by these populations. Thus, few theorists would wish to attribute perfect efficacy to inhibition, lest it should fail to explain what it was enlisted to explain. So, the probability that inhibition will successfully deactivate the offending representation is less than 1.0 as a rule, and increasingly more so for populations thought to have deficits in inhibition.

If inhibition is imperfect, we must consider how inhibition failures should influence behavioral markers of inhibition. An *inhibition failure* occurs when inhibition does not return the intrusive item's activation to a point either at or below its baseline level by the time the attempt at inhibition ends. Although many factors may in principle contribute to inhibition failure, the degree of interference caused by a competitor surely contributes. Competition is a graded function that is related to the activation that a competitor possesses. Given imperfect inhibition, there will reach a point at which a competitor's activation level exceeds what inhibition can counter within the time interval of a trial, and this activation point should vary depending on the inhibition rate possessed by an individual. Regardless of this variation, however, inhibition theories predict that *the inhibition failure rate should be an increasing function of the degree of interference.*

What happens when inhibition fails? The most straightforward possibility is that the inhibition target may retain the activation level it possessed at the end of an inhibition attempt, which will be above baseline. If so, the item may persist in its activated state, enjoying facilitation it would not otherwise have accrued had the competitive incident not taken place. We shall refer to this as the *carryover assumption*. Inhibition theories thus predict that as interference increases, the proportion of trials on which there will be no inhibition, or even facilitation, ought to increase. These failures ought to be reflected in the behavioral index of inhibition. Thus, theories with imperfect inhibition imply a trade-off between the increasing demand for inhibition and the increasing likelihood of failure. Inhibition theories—regardless of whether they concern inhibition of memories, task sets, motor responses, or perceptual representations—predict the occurrence of a demand-success trade-off that should influence behavioral indices of inhibitory aftereffects. *The demand-success trade-off refers to the predicted tendency for behavioral indices sensitive to inhibition to follow a nonmonotonic function relating interference to impairment (increasing, and then decreasing), reflecting the joint influence of inhibition demand and failure rate.*

To see the impact of a demand-success trade-off, consider an example, using RIF as a model case. Suppose we had a way of knowing the amount of interference that each competitor caused during the retrieval practice phase. Suppose further that we knew the quantitative relationship between the level of interference and (a) inhibition demand (e.g., probability that inhibition would be triggered) and (b) inhibition success. What would the relationship between competition and these factors look like? How would they combine to produce the aggregate RIF effect on the final test? Inhibition demand should increase monotonically with the level of interference, while inhibition success should decrease monotonically (see Figure 6.5a and b).

The carryover assumption described earlier about the consequences of inhibition failure yields key predictions about observed inhibitory aftereffects. In general, observed inhibition will increase as interference increases, owing to the increased probability of triggering inhibition; however, as interference grows further, the failure rate may increase enough to counterbalance the increased likelihood of triggering inhibition. Indeed, as increasing interference causes the proportion of failed inhibition trials to grow, observed inhibition should decline and, ultimately, with high enough levels of failure, should turn to facilitation. Thus, inhibition models with imperfect inhibition and a carryover assumption predict a nonmonotonic change in inhibitory aftereffects as levels of interference are parametrically increased, increasingly so the more imperfect inhibition is (see Figure 6.5c). Note that, depending on the degree of facilitation enjoyed by the competing item, this behavioral facilitation effect may arise even when some fraction of interfering items are truly inhibited. This inhibition is masked by the abundance of facilitated competitors.

As this example illustrates, the moment that an inhibition model introduces imperfect inhibition, demand-success trade-offs will complicate the relationship between the amount of inhibition demanded and the amount that is ultimately measured. These dynamics can lead to counterintuitive findings that might, on their face, seem to contradict inhibition theory. For instance, if competitors are especially interfering, retrieval practice may cause no RIF or even facilitation of items that one might expect to be inhibited. Similarly, if a study manipulates the degree of interference (e.g., low vs. high) to see whether inhibition is greater in the high-interference condition, a variety of results might occur, depending on where the low- and high-interference points are on the interference continuum. If they are at low and middle points, greater inhibitory aftereffects will be observed for higher interference; if they are at middle and high points, the reverse will be observed. Similarly, the function relating interference to

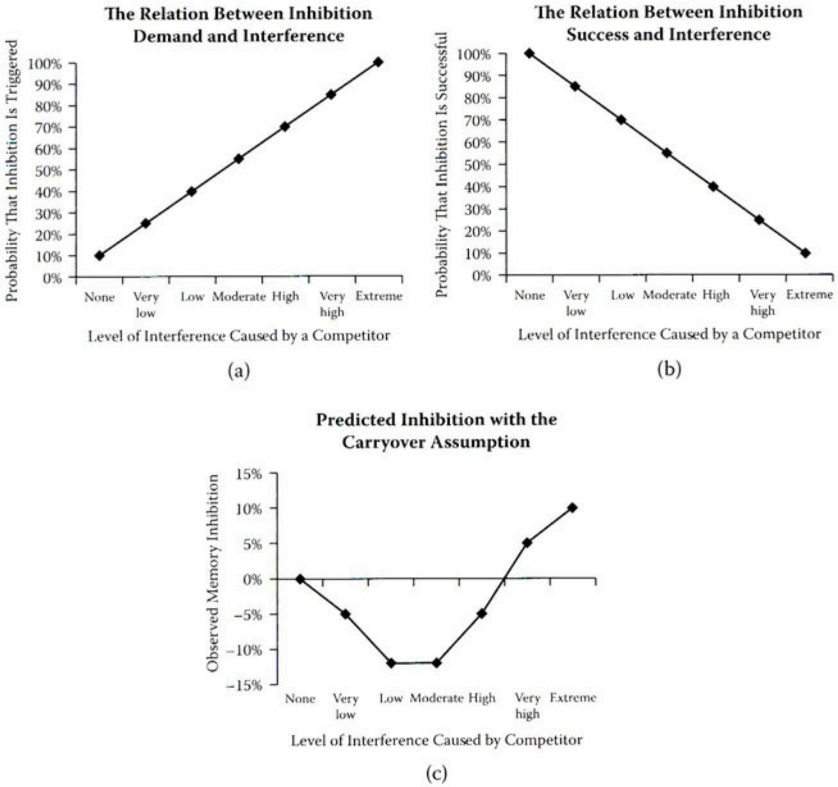


**Figure 6.5** The relationship between interference and inhibition. (a) As the interference caused by competitors increases, the demand for inhibition rises. (b) As the interference caused by competitors increases, so too does the likelihood that inhibition will fail to overcome the unwanted activation of the competitor. (c) The integration of the two functions predicted in a and b, under the carryover assumption. Here inhibition failure results in facilitation of highly interfering competitors, reversing the inhibition effect.

inhibition aftereffects will vary with population differences in inhibitory control. For low inhibitory control populations, the relationship between inhibition and interference ought to be more nonmonotonic than it is for high-functioning individuals, for whom inhibition failures may make up a tiny fraction of inhibition attempts. Thus, in one group (e.g., young subjects), more inhibition may be found for high- than for low-interference competitors, but for another group (e.g., older subjects), lower in inhibitory function, the opposite may be true, as the high-strength competitor will, for this group, instigate more failures. Thus, demand-success tradeoffs must be considered carefully in theoretical tests of inhibition, and also in testing inhibition deficit hypotheses.

### *Examples of Demand-Success Trade-Offs in Memory Inhibition*

The importance of demand-success trade-offs is best illustrated by considering the plight of a hypothetical investigator who decides to conduct research on inhibition. Suppose that our protagonist knows that interference triggers inhibition, and so ensures that the to-be-inhibited item is highly interfering. To begin, the investigator decides to simply look at the effects of a single inhibition attempt on later memory, compared to a baseline condition in which no inhibition is engaged. Being thorough, the investigator uses several paradigms, including semantic, episodic, and phonological RIF, as well as the think/no-think procedure.

Our investigator would doubtless be disappointed if he or she obtained the results depicted in Figure 6.6a. In every paradigm, performing a single suppression attempt on the highly competitive item facilitated its recall on a later test. Such a consistent pattern would lead our investigator to have grave doubts about whether inhibition combats interference. Not giving up, however, suppose that he or she persists, believing that a single inhibition attempt might not have been enough to inhibit competitors. So he or she includes 16 inhibition attempts, ensuring that people have ample opportunity to inhibit the unwanted memory. Surprisingly, he or she observes that the to-be-inhibited items show little sign of being inhibited below baseline, with the exception of a single study (see Figure 6.6a). Although the facilitation effects observed in the first series of studies have disappeared, and one study shows evidence of inhibition, this collection of eight experiments would justifiably be viewed as discouraging, and lead one to suspect that inhibition may not be involved in overcoming interference after all.

The data represented in Figure 6.6a are not fictitious, but were collected in our laboratory (Johnson & Anderson, 2004; Levy & Anderson, 2001; Levy et al., 2007; Shivde & Anderson, 2001). The main difference from our hypothetical case, however, was that the differing numbers of suppression attempts were not manipulated across experiments. Rather, they represented two of four points in within-subjects parametric manipulations of the number of suppression attempts. The full range of these parametric manipulations is depicted in Figure 6.6b.

Several unifying features may be observed in these otherwise diverse studies. First, in each study, competition was potent, so inhibition tended to initially fail and produce facilitation rather than below-baseline performance. These facilitation effects are most readily observed on a single suppression attempt, on which the participant is least prepared to combat the challenging interference posed by the newly encountered item. So, for example, after attempting to generate an associate to the

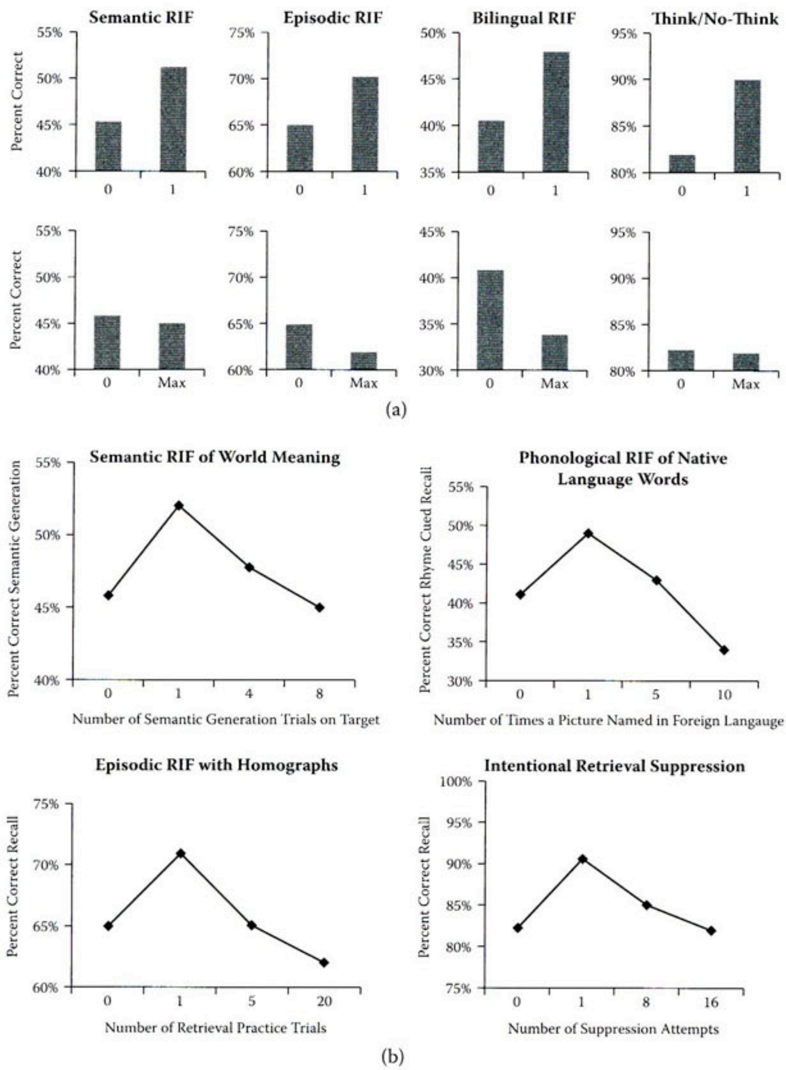


**Figure 6.6** Illustration of the two-point problem. (a) Illustrations of typical two-point comparisons across four paradigms designed to test for inhibition. The top row contrasts a baseline condition with a single suppression attempt in four studies; none show inhibition, and all show reliable facilitation. The second row illustrates a comparison between the baseline condition and a higher number of suppression attempts. Overall, the pattern does not provide support for inhibitory aftereffects (except in one study). (b) Parametric manipulations in these four paradigms reveal that all four involve an initial inhibition failure that boosts target recall, but which is followed by successful engagement of inhibition. Two-point assessments hide these nonmonotonic functions, which clearly indicates the involvement of inhibition in all cases.

verb meaning of *prune* in response to prune *t _ _ m*, participants often reported involuntarily thinking, at first, of the inappropriate noun meaning of *prune*; as a result, they were later more likely to successfully generate *fruit* on a semantic generation test than they would have been had they not attempted to practice the *trim* meaning of *prune* —in apparent contradiction to the idea that inhibition was recruited to suppress the irrelevant noun meaning. This initial rise in recall indicates that people failed to inhibit the competitor, and that they had difficulty with the retrieval practice task itself, consistent with the very low retrieval practice success often observed on the first trial in this design.

The second common feature of these data is that after the initial failure to inhibit, further inhibition attempts reduce performance monotonically. In some cases, repeated suppression impairs the to-be-inhibited item below baseline, whereas in other cases it does not. In every case, however, there is a significant decline in performance across repetition levels. Because one inhibition attempt yields significant enhancement of the to-be-inhibited item, and we know that this facilitation persists through to the final test, the decline in performance with increasing repetitions implies an active process that counters that elevated accessibility, even if performance at the highest level fails to go below baseline. The failure to go below baseline may, in some cases, reflect the aggregate influence of items that persistently intrude and remain facilitated, and others that are progressively inhibited below baseline; alternatively, all items may be initially facilitated, and the decline from peak recall may reflect increasing success at countering their intrusion. In either case, the rise in recall after a single suppression argues in favor of the carryover assumption. Interestingly, despite the carryover, which ought to make the competitor even more intrusive, people appear able to dynamically adjust cognitive control to better handle high levels of competition on subsequent trials—a finding reminiscent of evidence for conflict adaptation in studies of executive control (see Botvinick, Cohen, & Carter, 2004). These findings are representative of what happens in a variety of inhibition paradigms when interference is particularly high. Instead of finding greater below-baseline inhibition, one gets null effects or even facilitation, consistent with a demand-success trade-off.

The findings provide an important cautionary tale for investigators interested in studying inhibition. Surveying the literature on inhibitory processes, whether in memory or other domains, reveals that the overwhelming majority of experiments employs only two measurement points: a baseline condition, where inhibition should be absent, and an experimental condition, where it is believed to be present. The problem with using only two measurement points—which we refer to as the *two-*

*point problem*—is vividly illustrated by the plight of our hypothetical investigator. This investigator might seem to have done everything right, yet ended up with no evidence for inhibition, even after multiple suppression attempts. This investigator's problem arises because he chose competitors that generated maximal interference, without appreciating the significance of demand-success trade-offs, and because he only measured inhibition with two points. Had he included a parametric manipulation of the number of inhibition attempts, the experimenter would not have been misled about the role of inhibition in these tasks.

These nonmonotonic functions also suggest that demand-success trade-offs have strong potential to mislead investigators interested in theoretical models of inhibition and inhibitory deficits. Consider again a hypothetical investigator who manipulates the level of interference to evaluate whether inhibition is involved in a task. The two-point data (Figure 6.7a) might indicate less impairment in the high-interference condition, appearing to challenge an inhibition account. This isn't necessarily true, though, as there may be an initial failure that is followed by gradual inhibition (Figure 6.7b). Thus, a parametric manipulation may have provided a crucial window into the dynamics of inhibition.

Demand-success trade-offs are also relevant to research on inhibitory deficits. Consider the results of a hypothetical experiment looking at the effects of aging on memory inhibition (Figure 6.7d). These data indicate less memory inhibition in older adults, when comparing a single experimental condition to a baseline. Here again, this may not be true. We cannot tell from two measurement points whether older adults are truly deficient at inhibition (Figure 6.7e), or if they have difficulty engaging control on the first trial, followed by highly effective suppression thereafter (Figure 6.7f). Indeed, in the latter case, even if recall never goes below baseline, there is actually more inhibition, when compared to the initial rise. Because the initial failure elevates the accessibility of the to-be-inhibited item and that persists, one must account for this rise in interpreting the data in all other repetition conditions. One credible account, for example, is that older adults are somewhat slower in recruiting inhibition (or any other process, see, e.g., Salthouse, 1996) at first (leading to an initial failure, and the rise in performance), but once inhibition has been engaged, it is as effective as it is in younger populations.

The foregoing examples illustrate the power of demand-success trade-offs to mislead investigators interested in inhibition. Despite substantial interest in inhibitory deficits, many investigators have not considered these trade-offs. To the extent that inhibitory control has been assessed with experimental designs that neglect these issues, as we have argued, the literature should be plagued by inconsistencies in

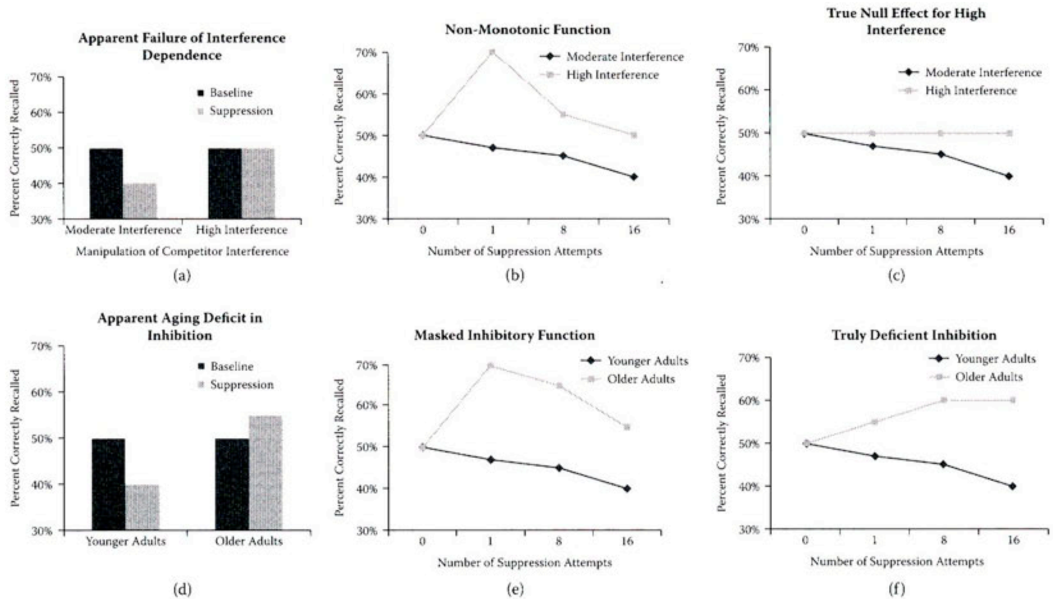


**Figure 6.7** Examples of how the demand-success trade-off influences assessment of inhibition theories. (a) Hypothetical results from an investigator who manipulates interference to evaluate whether inhibition is involved in a task. The data indicate less impairment (e.g., RIF) in the high-interference condition, appearing not to support an inhibition account. However, we cannot tell, from two measurement points, whether (b) the underlying function is an initial failure for highly interfering items followed by successful engagement of inhibitory control, or (c) there really is no inhibition in the high-interference condition. (d) A similar situation may arise in testing for variation in inhibitory control abilities across different populations. Here hypothetical data suggest that the older group has an inhibitory deficit, compared to the younger group. From this two-point assessment, one cannot be sure whether older adults (e) experience an initial inhibition failure followed by successful engagement of inhibition or (f) are truly deficient at inhibitory control.

the support of inhibitory deficits, generating reasonable doubt about their utility. We argue that such variability is not necessarily a sign for or against these theories, but rather reflects the insensitivity of these measures for assessing inhibition. As a solution, we argue for the use of parametric manipulations. By tracking the development of inhibition with repeated application, one can distinguish between truly null inhibitory effects and trade-offs arising from initial failures to engage inhibition. We believe that the logic underlying this method can be adapted to any domain concerned with inhibition.

## CONCLUDING REMARKS

With rare exception, the history of research on memory has adopted the point of view that forgetting is a problem to be overcome. This simple and pervasive assumption underlies nearly a century of research on forgetting that has emphasized the contributions of passive mechanisms, such as decay, interference, and contextual fluctuation. One of the seminal contributions of Robert Bjork has been to question this foundational premise, both empirically and theoretically. Empirically, in a career spanning four decades, Bjork introduced experimental methods such as the item and list method directed forgetting paradigms to establish that people can, in fact, forget things on purpose, and that it is useful to do so. Theoretically, Bob and Elizabeth have been articulate champions of the idea that forgetting, far from being a problem, can be an adaptive feature of the cognitive system and that this function may be accomplished by inhibition. These ideas, fittingly developed in a festschrift volume in honor of Endel Tulving (Bjork, 1989), inspired the first author to develop work on RIF. Two decades later, the fruits of this collaboration with Bob and Elizabeth Bjork can be seen in the sizable volume of work on memory inhibition and in the acceptance of the view that forgetting can indeed be functional.

In this chapter, we have examined evidence for a basic assumption of this adaptive forgetting view—that inhibition mechanisms are engaged specifically to counter unwanted accessibility. According to this view, when persisting accessibility of a trace hinders some cognitive process—whether it is retrieval or attention to other representations—inhibition is triggered to overcome this interference. Over the last two decades, considerable evidence has amassed for this assumption, and we reviewed this evidence, focusing, in particular, on work with the RIF and think/no-think paradigms. Evidence for interference-dependence has been obtained by manipulating the competitive dynamics during retrieval practice, either by varying the a priori associative strengths of competitors or by manipulating the demands imposed on cognitive control. Collectively, the recurring pattern indicates that when the task goal can be disrupted by heightened accessibility of a nontarget memory, people are more likely to show forgetting of the interfering trace on later tests, riffing (pun intended) nicely on Bjorkian themes of adaptive forgetting.

Although the evidence clearly favors the interference-dependence assumption, increasing the accessibility of an unwanted memory will not always lead to increases in memory inhibition. In the last section, we argued that unless one assumes a perfectly effective inhibition process, behavioral markers of the aftereffects of inhibition should not uniformly increase as the interference a competitor causes increases. Rather, when inhibition is imperfect, the size of the inhibitory aftereffect should follow a nonmonotonic function, whereby it initially increases with interference, and then decreases as interference exceeds the person's capabilities. We illustrated these trade-offs, and how they can be problematic. These trade-offs, largely neglected, have the potential to undermine theoretical inferences in any domain concerned with inhibition. Indeed, failure to consider demand-success trade-offs contributes considerable noise to the empirical case for inhibition. We argued for the utility of parametric designs to detect demand-success trade-offs, and better quantify the involvement of inhibition.

Adaptive forgetting plays an important role in the regulation of memory that may be important to our broader functioning. Indeed, some have argued for its role in maintaining a healthy sense of perspective (Pronk, Karremans, Overbeek, Vermulst, & Wigboldus, 2010). For example, we are confident that Bob Bjork can recount, in glorious detail, the events of his triumphant season as coach of the UCLA Psychology Department's flag football team during the first author's time as his graduate student; this is true, even while we very much doubt Bob would recall very much about the many times during that same period that the first author pestered Bob relentlessly for comments on their first manuscript on RIF (success at which was negatively correlated with flag football events). Even after two decades of research on RIF and four on directed forgetting, we can think of no better proof for adaptive forgetting than this. The very best research ideas often come to us from our personal experience, and in this regard, it should perhaps come as no surprise to those who know Bob that he championed the virtues of forgetting so admirably. And the field, and his friends, are all the better for it.

## REFERENCES

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

Anderson, M. C. & Bjork, R. A. (1994). Mechanisms of inhibition in long-term memory: A new taxonomy. In D. Dagenbach & C. 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., & Green, C. (2001). Suppressing unwanted memories by executive control. *Nature, 410*, 131–134.

Anderson, M. C., & Levy, B. J. (2009). Suppressing unwanted memories. *Current Directions in Psychological Science, 18*, 184–194.

Anderson, M. C., Ochsner, K., Kuhl, B., Cooper, J., Robertson, E., Gabrieli, S. W., et al. (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.

Bajo, M. T., Gomez-Ariza, C. J., Fernandez, A., & Marful, A. (2006). Retrieval-induced forgetting in perceptually driven memory tests. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32*, 1185–1194.

Bäuml, K. (1996). Revisiting an old issue: Retroactive interference as a function of the degree of original and interpolated learning. *Psychonomic Bulletin and Review, 3*, 380–384.

Bäuml, K. (1997). The list-strength effect: Strength-dependent competition or suppression. *Psychonomic Bulletin and Review, 4*, 260–264.

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

Bjork, R. A. (1970). Positive forgetting: The noninterference of items intentionally forgotten. *Journal of Verbal Learning and Verbal Behavior, 9*, 255–268.

Bjork, R. A. (1975). Retrieval as a memory modifier: An interpretation of negative recency and related phenomena. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola Symposium* (pp. 123–144). Hillsdale, NJ: Erlbaum.

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

Botvinick, M. M., Cohen, J. D., & Carter, C. S. (2004). Conflict monitoring and anterior cingulate cortex: An update. *Trends in Cognitive Sciences, 8*, 539–546.

Campbell, J. I. D., & Phenix, T. L. (2009). Target strength and retrieval-induced forgetting in semantic recall. *Memory and Cognition, 37*, 65–72.

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

Dagenbach, D., & Carr, T. H. (Eds.). (1994). *Inhibitory processes in attention, memory, and language*. San Diego, CA: Academic Press.

Dagenbach, D., Carr, T. H., & Barnhardt, T. M. (1990). Inhibitory semantic priming of lexical decisions due to failure to retrieve weakly activated codes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*, 328–340.

Dempster, F. N., & Brainerd, C. J. (Eds.) (1995). *Interference and inhibition in cognition*. San Diego, CA: Academic Press.

Depue, B. E., Banich, M. T., & Curran, T. (2006). Suppression of emotional and non-emotional content in memory: Effects of repetition on cognitive control. *Psychological Science, 17*, 441–447.

Depue, B. E., Curran, T., & Banich, M. T. (2007). Prefrontal regions orchestrate suppression of emotional memories via a two-phase process. *Science, 37*, 215–219.

Dolan, R. J., & Fletcher, P. C. (1997). Dissociating prefrontal and hippocampal function in episodic memory encoding. *Nature, 388*, 582–585.

Fletcher, P. C., Shallice, T., & Dolan, R. J. (2000). "Sculpting the response space"—an account of left prefrontal activation at encoding. *Neuroimage, 12*, 404–417.

Gorfein, D. S., & MacLeod, C. M. (Eds.). (2007). *Inhibition in cognition*. Washington, DC: American Psychological Association.

Henson, R. N., Shallice, T., Josephs, O., & Dolan, R. J. (2002). Functional magnetic resonance imaging of proactive interference during spoken cued recall. *Neuroimage, 17*, 543–558.

Johansson, M., Aslan, A., Bauml, K. H., Gabel, A., & Mecklinger, A. (2007). When remembering causes forgetting: Electrophysiological correlates of retrieval-induced forgetting. *Cerebral Cortex, 17*, 1335–1341.

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

Kuhl, B. A., Dudukovic, N. M., Kahn, I., & Wagner, A. D. (2007). Decreased demands on cognitive control following memory suppression reveal benefits of forgetting. *Nature Neuroscience, 10*, 908–914.

Levy, B. J., & Anderson, M. C. (2001, June). *Regulation of conscious awareness: Further evidence for a direct suppression mechanism*. Poster presented at the Thirteenth Annual Meeting of the American Psychological Society, Toronto, ON, Canada.

Levy, B. J., Kuhl, B. A., & Wagner, A. D. (2010). The functional neuroimaging of forgetting. In S. Della Sala (Ed.), *Forgetting*. Hove & New York: Psychology Press.

Levy, B. J., McVeigh, N. D., Marful, A., & Anderson, M. C. (2007). Inhibiting your native language: The role of retrieval-induced forgetting during second-language acquisition. *Psychological Science, 18*, 29–34.

Linck, J. A., Kroll, J. F., & Sunderman, G. (2009). Losing access to the native language while immersed in a second language: Evidence for the role of inhibition in second-language learning. *Psychological Science, 20*, 1507–1515.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience, 24*, 167–202.

Pronk, T. M., Karremans, J. C., Overbeek, G., Vermulst, A. A., & Wigboldus, D. H. J. (2010). What it takes to forgive: When and why executive function facilitates forgiveness. *Journal of Personality and Social Psychology, 98*, 119–131.

Salthouse, T. A. (1996). The processing-speed theory of adult age differences in cognition. *Psychological Review, 103*, 403–428.

Saunders, J., Fernandes, M., & Kosnes, L. (2009). Retrieval-induced forgetting and mental imagery. *Memory and Cognition, 37*, 819–828.

Shimamura, A. P., Jurica, P. J., Mangels, J. A., Gershberg, F. B., & Knight, R. T. (1995). Susceptibility to memory interference effects following frontal lobe damage: Findings from tests of paired-associate learning. *Journal of Cognitive Neuroscience, 7*, 144–152.

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, DC: American Psychological Association.

Storm, B. C., Bjork, E. L., & Bjork, R. A. (2007). When intended remembering leads to unintended forgetting. *Quarterly Journal of Experimental Psychology, 60*, 909–915.

Storm, B. C., Bjork, E. L., Bjork, R. A., & Nestojko, J.F. (2006). Is retrieval success a necessary condition for retrieval-induced forgetting? *Psychonomic Bulletin and Review, 13*, 1023–1027.

Storm, B. C., & Nestojko, J. F. (2009). Successful inhibition, unsuccessful retrieval: Manipulating time and success during retrieval practice. *Memory, 18*, 99–114.

Wimber, M., Rutschmann, R. M., Greenlee, M. W., & Bäuml, K. (2009). Retrieval from episodic memory: Neural mechanisms of interference resolution. *Journal of Cognitive Neuroscience, 21*, 538–549.