Available online at www.sciencedirect.com
ScienceDirect
Acta Psychologica 127 (2008) 623–635
acta psychologica
www.elsevier.com/locate/actpsy

Individual differences in the suppression of unwanted memories: The executive deficit hypothesis

Benjamin J. Levy *, Michael C. Anderson

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

Received 21 March 2007; received in revised form 29 November 2007; accepted 3 December 2007
Available online 1 February 2008

Abstract

When confronted with reminders to an unpleasant memory, people often try to prevent the unwanted memory from coming to mind. In this article, we review behavioral and neurocognitive evidence concerning the consequences of exerting such control over memory retrieval. This work indicates that suppressing retrieval is accomplished by control mechanisms that inhibit the unwanted memories, making them harder to recall later, even when desired. This process engages executive control mechanisms mediated by the lateral prefrontal cortex to terminate recollection-related activity in the hippocampus. Together, these findings specify a neurocognitive model of how memory control operates, suggesting that executive control may be an important means of down-regulating intrusive memories over time. We conclude by proposing that individual differences in the regulation of intrusive memories in the aftermath of trauma may be mediated by pre-existing differences in executive control ability. In support of this *executive deficit hypothesis,* we review the recent work indicating links between executive control ability and memory suppression.
Published by Elsevier B.V.

PsycINFO classification: 2340

Keywords: Executive control; Inhibition; Forgetting; Trauma; Memory; Suppression

1. Inhibiting unwanted memories: Individual differences and the executive deficit hypothesis

Over the last several years, humanity has witnessed some of the most horrific natural disasters that we are likely to see in our lifetime. The Asian tsunami and the Kashmir earthquake each killed hundreds of thousands of people. Hurricane Katrina caused a flood that engulfed an American city, displacing its residents who lost everything they owned, and whose lives were fundamentally altered. The impact of such events is not limited to the loss of physical property and livelihood. Such experiences have a profound psychological impact, with people undergoing a long period of mental adaptation and adjustment. Part of this adaptation is the process of learning to live with intrusive remindings that disrupt the ability to cope. When simple objects or events remind a person of their trauma, people frequently report trying to exclude the unwanted experience from mind, so that they can regain their focus and go about their daily activities. Although the frequency of these remindings diminish over time for most people, there are striking individual differences in the rate at which this adjustment occurs; for some, intrusions diminish rapidly as the person adjusts, whereas for others intrusions may continue for years or, in the extreme, decades, significantly disrupting a person's mental and physical well-being. These observations raise several questions. What are the cognitive and neural mechanisms by which people try to control memory? What effects do such efforts have on the memories themselves? Why do people vary so dramatically in this ability? Answering these questions will prove fundamental to the ability of psychological science to ease the suffering of individuals in the aftermath of a traumatic experience.

* Corresponding author.
E-mail address: blevy@uoregon.edu (B.J. Levy).

0001-6918/$ - see front matter Published by Elsevier B.V.
doi:10.1016/j.actpsy.2007.12.004

In this article, we review a program of research that addresses these questions. A core claim is that people control the influence of unwanted memories by engaging executive control mechanisms that target the neural regions involved in declarative memory. Engaging these control processes suppresses the unwanted memory, rendering it more difficult to retrieve in the future, even when it is desired. We first review work documenting this ability, isolating the inhibitory control mechanisms that underlie it, and establishing a neurocognitive model of how such mechanisms take place. We then argue that whereas studying the general functioning of these mechanisms is vital, it is equally important to understand individual differences and the factors that produce this variability. A key claim, which we term the *executive deficit hypothesis,* is that differences in the ability to combat intrusive memories arise, in part, from differences in executive control abilities. We then review work on sources of individual differences.

1.1. Overview of work on memory control

An overarching framework guiding our research on controlling unwanted memories is that this ability is analogous to controlling overt behavior. It is clear that once started, motor actions can be stopped. Imagine a basketball player initiating the motion of shooting a three-pointer. If, at this point, a defender arrives and threatens his shot, he can override this motor movement and prevent his shot from being blocked, either by altering his shot to provide more arc or by simply stopping his shooting motion. It is our proposal that this ability to control overt behavior is also what allows us to control internal thought. Both cases are instances of a general 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, weaker response to that stimulus is desired. Without this ability, we would be slavishly limited to our first instinct, our habitual responses. In many instances, our desire is to perform a novel action or to think of something else, and, in fact, this ability to flexibly adapt behavior is widely regarded as a hallmark of intelligence.

The specific proposal advanced here and elsewhere is that we accomplish this control through inhibition of the prepotent response, both when that response is an overt action or a memory. 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 is engaged to weaken the dominant memory, enabling selective control over what is retrieved or simply stopping retrieval altogether. We argue that this action has a lasting effect, leading to later memory impairments for the avoided memories even when we want to retrieve them. In the following section, we discuss evidence that uniquely supports this inhibitory control perspective in both response override situations mentioned earlier: the need to selectively retrieve a memory and the desire to stop retrieval.

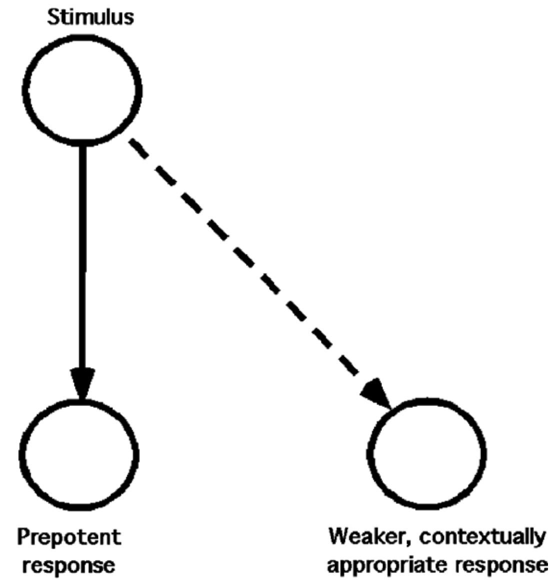

Fig. 1. A typical response-override situation. In this figure, a stimulus is associated with two responses, one of which is stronger (prepotent), and the other of which is weaker (indicated by the dotted line). Response override occurs whenever one needs to either select the weaker, but more contextually 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.

1.1.1. Selective retrieval

When attempting to retrieve a specific memory, it is rarely the case that the desired memory is the only memory related to the cues that guide retrieval. Rather, retrieval cues are frequently related to a broad range of memories, many of which are much more strongly related to the cues than the memory that one is currently attempting to retrieve. In these situations, the non-target memories compete for access to conscious awareness, necessitating some process to allow selective retrieval of the non-dominant memory. For example, when reminiscing about the name of a childhood friend's father, memories of childhood friends, one's own father, or similar-sounding names all threaten to overwhelm the rememberer and thwart the attempt to retrieve that specific name. According to our framework, this type of selective retrieval represents a paradigmatic case of response override, in which one must select a weaker memory in the face of interference from one or more prepotent competitors. If inhibition is engaged to stop prepotent memories from coming to mind (thus promoting selective retrieval), perhaps inhibition would induce memory impairment for the competitors. Thus, the very act of remembering might cause forgetting of related memories.

This prediction has been explored in a procedure known as the retrieval practice paradigm (Anderson, Bjork, & Bjork, 1994). In a typical version of this paradigm, subjects study lists of category-exemplar words pairs (e.g., Fruits-Orange, Fruits-Banana, Drinks-Scotch). They then practice a subset of these items by 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. As one might expect, practiced items (Orange) are recalled more often than items from unpracticed categories (Scotch), which act as a baseline measure of how well items are recalled when that category is not practiced (see Fig. 2). More interestingly, the unpracticed items from the practiced categories (Banana) are recalled less often than the baseline items. Thus, retrieving some items during the practice phase 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). In fact, RIF applies to a strikingly broad range of situations. For instance, when students practice retrieving facts about a topic they have just studied, they have a harder time retrieving unpracticed facts about that topic on short answer or essay questions (Carroll, Campbell-Ratcliffe, Murnane, & Perfect, 2007). Similarly, when a speaker retrieves a second-language word to name an object, they later have difficulty generating the phonology of the corresponding native-language word (Levy, McVeigh, Marful, & Anderson, 2007). Simply generating examples of a category from semantic memory causes later impairments for dominant, unpracticed memories from those categories, even if those competitors were not studied earlier (Johnson & Anderson, 2004; for a related finding see Starns & Hicks, 2004). The generality of this empirical finding suggests that RIF is a general phenomenon that impairs memory whenever unwanted items intrude during retrieval.

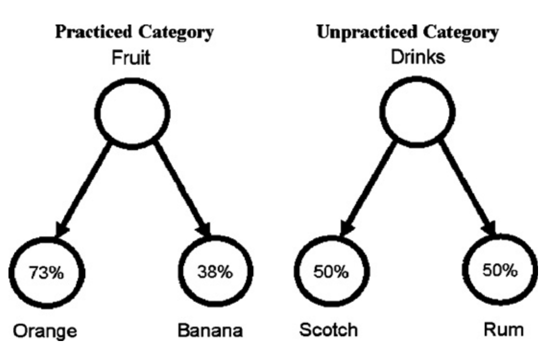

Fig. 2. A standard categorical retrieval-induced forgetting study. Illustrated here are two items from each of two categories that subjects have studied (typically six items are studied from eight categories). In this example, subjects perform retrieval practice on "Fruit-Orange," but not on "Fruit-Banana" (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, retrieval 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 on items from baseline categories (Scotch and Rum).

Despite our emphasis on inhibition, the basic finding of RIF is compatible with several non-inhibitory mechanisms. One such explanation is that practiced items are 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., Anderson, 1983; Raaijmakers & Shiffrin, 1981), in which 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 instance, subjects may simply unlearn the association between "Fruits" and "Banana" or bias the representation of that category toward the practiced items (retrieving "Orange" and "Lemon" might make you interpret the cue "Fruits" as being "Citrus Fruits"). Critically, none of these alternative explanations claim that any change is occurring to the item itself (see Anderson & Bjork, 1994, for a review of these non-inhibitory sources of impairment). A core claim of the inhibitory control perspective, on the other hand, is that there is a decrease in the activation level of the unwanted item *itself*.

Several properties of RIF uniquely support the involvement of inhibition. First, if the non-retrieved responses are truly inhibited, then RIF should be observed regardless of what cue is used to test the memory. In other words, forgetting should be *cue-independent* and, therefore, generalize to novel retrieval cues in the test phase, rather than being specific to the cues 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 new, independent cue (e.g., 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. Associative blocking, for example, cannot explain this type of impairment because the new independent cue is unrelated to the strengthened exemplars; thus there is no reason why they should block retrieval ("Monkey" is unrelated to "Orange", so there is no reason that "Orange" would block retrieval of "Banana"). Cue-independent forgetting has now been observed many times (e.g., Anderson & Bell, 2001; Anderson, Green, & McCulloch, 2000; Anderson & Spellman, 1995; Aslan, Bäuml, & Pastötter, 2007; Camp, Pecher, & Schmidt, 2005; Johnson & Anderson, 2004; Levy et al., 2007; MacLeod & Saunders, 2005; Radvansky, 1999; Saunders & MacLeod, 2006), 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; Verde, 2004) and implicit lexical decision tests (Veling & van Knippenberg, 2004).

Another difference between the inhibitory and non-inhibitory accounts of RIF concerns the relationship between strengthening competitors and memory impairment. Non-inhibitory explanations of RIF, such as associative interference, predict that memory impairment should occur whenever competitors are strengthened. Several findings, however, have shown that strengthening practiced items, by itself, does not impair recall of the competitors. For example, repeated study exposures to the practiced items strengthen later retrieval of those word pairs to a similar degree as does retrieval practice, yet this type of strengthening does not impair competing items, provided that output interference at test is controlled[1] (Anderson, Bjork, & Bjork, 2000; Bäuml, 1996, 1997, 2002; Ciranni & Shimamura, 1999). If strengthening can occur without impairing related items, it suggests that retrieval practice, not strengthening, 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 referred to as *interference-dependence*. Anderson et al. (1994) found that retrieval practice causes significantly more RIF for high frequency competitors (e.g., Orange) than it does for low frequency competitors (e.g., Kiwi), consistent with the idea that only interfering items are inhibited. The greater RIF for high frequency members is observed regardless of whether one performs retrieval practice on high or low frequency exemplars, and is independent of how much those practiced items are strengthened. Similarly, retrieving the subordinate meaning of an asymmetric homograph significantly impairs access to the dominant meaning, but the converse is not true (e.g., Shivde & Anderson, 2001). This poses a problem for non-inhibitory accounts of RIF: 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. Strikingly, it appears that successful retrieval practice of target items is not even necessary to observe RIF of competing items, as RIF occurs even when strengthening of practiced items is not possible (Storm, Bjork, Bjork, & Nestojko, 2006). Each of the aforementioned properties of RIF strongly implicates inhibition, while providing serious challenges for any non-inhibitory explanation. 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.

1.1.2. Stopping retrieval

The need to control memory is not limited to selecting a non-dominant response; it is also sometimes necessary to simply prevent a memory from coming to mind. When presented with a reminder of something upsetting (e.g., glimpsing a photograph of a loved one who has recently died), one often desires to put such unpleasant thoughts out of mind. Can inhibitory mechanisms be engaged to serve these goals as well? To address this question, Anderson and Green (2001) developed the "think/no-think" (TNT) paradigm by adapting the "go/no-go" task, which has been widely used to study the suppression of motor responses in both humans (e.g., de Zubicaray, 2000; Garavan, Ross, & Stein, 1999) and monkeys (Sakagami & Niki, 1994; Sasaki, Gemba, & Tsujimoto, 1989). In a typical "go/no-go" task, subjects must provide a motor response (e.g. pressing a button) to a variety of stimuli (e.g. letters). However, a specific, infrequently presented stimulus (e.g. the letter X) requires the suppression of that response. Subjects' ability to withhold the response is taken as a metric of effective inhibitory control.

In the TNT paradigm, participants learn a list of cue-target word pairs (e.g., Ordeal–Roach), rather than motor responses. Then subjects are presented with the cue they studied earlier (Ordeal) and are either asked to think of the associated word (Roach) or to prevent that associated word from coming to mind, just as "go/no-go" subjects had to prevent motor responses from occurring. Unlike the "go/no-go" task, however, it is not possible to observe whether or not a person retrieves a memory. Instead, Anderson and Green looked at the subsequent accessibility of the response words after having seen the cues as many as 16 times during the TNT phase. If subjects are able to recruit inhibitory control mechanisms to suppress the unwanted memories on "no-think" trials and this suppression lingers, then we should expect these words to be less accessible later. To assess this, subjects were again presented with the same cues they studied earlier (Ordeal) and asked to provide the correct target memory (Roach) for all of the previously studied word pairs.

As shown in Fig. 3, there is a large difference in the final recall performance between the "think" and "no-think" conditions, reflecting the total *memory control effect*, of which *positive control* (facilitation of "think" items above

[1] To control for output interference, retrieval-induced forgetting studies often use letter stems (e.g., fruits-o_____) to specify the order in which targets should be retrieved during the final test, allowing them to ensure that the non-practiced items are tested before the practiced items. Even when controlling for output interference in this way, retrieval-induced forgetting is observed (Anderson et al. (1994); Anderson & McCulloch, 1999; Anderson et al., 2000; Anderson & Bell, 2001; Bäuml, 2002; Bäuml & Hartinger, 2002). If these precautions are not taken and subjects are allowed to recall items in an unconstrained fashion (e.g., free recall or category cued recall without constraining letter stems) then subjects will tend to recall practiced items first, thus producing output interference for unpracticed items from those categories. Critically, while this will lead to forgetting of unpracticed competitors in the extra exposure condition, this impairment is produced entirely during the test rather, not by the strengthening phase where the extra exposures occurred.

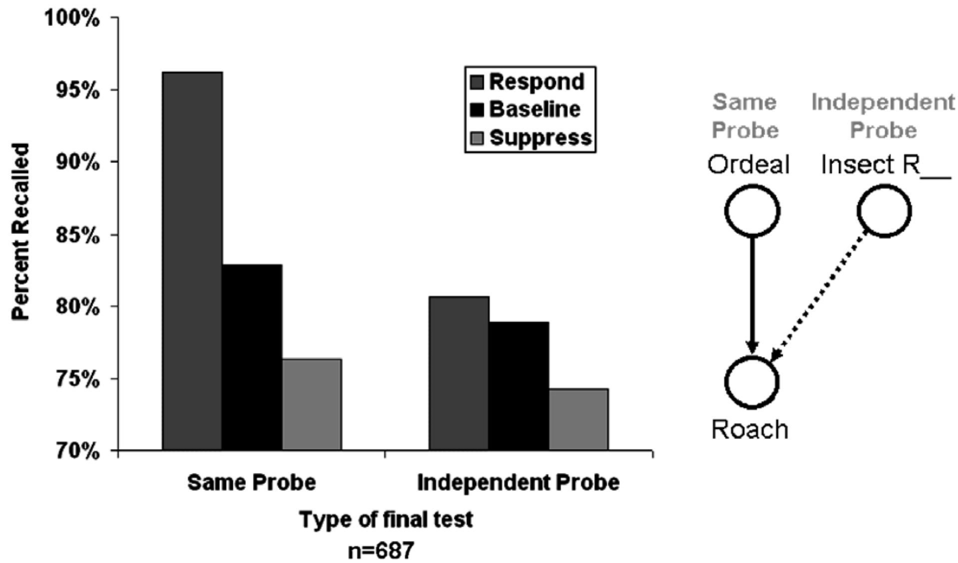

Fig. 3. Final recall performance in the think/no-think (TNT) procedure. The graph shows the percentage of items that subjects correctly recalled on the final test as a function of whether they tried to recall the item (Respond), suppressed the item (Suppress), or had no reminders to the item during the TNT 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).

baseline due to subject-initiated retrieval) and *negative control* (inhibition of "no-think" items below baseline) are components. Unsurprisingly, when inclined to remember (i.e., the "think" condition), reminders enhance memory; but when people desire not to be reminded ("no-think" condition), the reminders not only fail to enhance memory, they set the occasion for processes that impair memory. The strongest evidence for inhibition comes from the *negative control effect*: subjects have more difficulty recalling items that they previously avoided thinking about than for baseline items. This finding is counterintuitive because one would expect that repeated reminders to an item should make it more accessible, not less accessible.[2] The data presented in Fig. 3 come from a meta-analysis of 687 subjects run in the TNT procedure in our laboratory, providing clear evidence that the negative control effect is robust across a large sample.[3] In addition, the *negative control effect* has now been replicated multiple times (Anderson et al., 2004; Bergstorm, Richardson-Klavehn, & de Fockert, 2006; Depue, Banich, & Curran, 2006; Hertel & Calcaterra, 2005; Hertel & Gerstle, 2003; Joorman, Hertel, Brozovich, & Gotlib, 2005; Wessel, Wetzels, Jelicic, & Merckelbach, 2005; Hotta & Kawaguchi, 2006; Anderson & Kuhl, in preparation; Anderson, Reinholz, Kuhl, & Mayr, in preparation; Bell & Anderson, in preparation; Depue, Curran, & Banich, 2007; Hart, 2006; although, see Bulevich, Roediger, Balota, & Butler, 2006, who observed a significant memory control effect, but not a negative control effect). Thus, it seems clear that people have the ability to prevent unwanted memories from coming to mind, and this results in the avoided memories being harder to recall later.

While the *negative control effect* appears to support the involvement of inhibitory processes in stopping retrieval, it is also compatible with non-inhibitory mechanisms. For example, people could generate diversionary thoughts whenever they see one of the "no-think" cues. Then, when presented with the studied cues during the final memory test (Ordeal), the diversionary thoughts may come to mind readily and block retrieval of the original response word. Critically, this explanation does not claim that the response word itself has been inhibited, simply that alternative associations interfere with the ability to access the target memory.

To address this possibility, Anderson and Green (2001) demonstrated that memory impairment is observed even when subjects are provided with novel, extralist items as retrieval cues on the final test (e.g., "Insect-R______" for "Roach"). As with RIF, these independent probes provide critical evidence that inhibition is the mechanism that produces the *negative control effect*. In support of the inhibitory control hypothesis, Anderson and Green (2001) also ruled out several alternative explanations of the observed memory impairment. They replicated the *negative control effect* when subjects were paid for correct responses and when subjects were misled about the experimenters' expectations immediately before the final test (by telling them

[2] Even if below-baseline inhibition did not occur, repeatedly avoiding retrieval when reminders are present clearly deprives a memory of reactivations (either due to intentional retrieval or to spontaneous reminding) that otherwise would preserve and enhance more desirable traces. Thus, even a failure to suppress below baseline suggests that subjects can effectively control memory.

[3] This meta-analysis contains all subjects who participated in our laboratory in the TNT procedures reported by Anderson and Green (2001); Anderson et al. (2004) between November 1999 and April 2004. This difference is, of course, highly significant (a simple *t*-test between the baseline and suppression items yields a *p*-value less than $.01 \times 10^{-13}$).

that attempting to suppress an idea would ironically make it more accessible). These findings rule out the possibility that subjects were simply withholding responses that they knew (e.g., due to demand characteristics). In their final experiment, Anderson and Green showed that simply asking subjects to avoid saying the response word eliminated the *negative control effect*, isolating the attempt to prevent the unwanted memory from coming into consciousness as being the critical factor for producing memory impairment. Taken together, these results indicate that the memory impairment arises from an inhibitory control mechanism that prevents the unwanted memory from entering awareness.

As we have argued elsewhere (Anderson, 2006; Anderson & Green, 2001; Anderson & Levy, 2002, 2006; Anderson et al., 2004), this body of work represents a viable model for studying the voluntary form of repression (suppression) proposed by Freud (1915/1963). Freud (1915/1963) defined repression as "simply the function of rejecting and keeping something out of consciousness" (p. 147), which is precisely what we asked our subjects to do. Given the evidence presented in this section, it seems clear that under certain circumstances, at least, voluntary repression can occur. However, we have also tried to be clear that the basic results described above cannot yet be fully extended to the type of traumatic forgetting often explained by Freudian repression. To do this, research with this laboratory model will need to steadily build the ecological validity necessary to make such claims (Anderson & Levy, 2006).

One critical difference between this line of research and memory suppression in naturalistic settings is that in actual situations, it is likely that the avoided thoughts will be emotional, negatively valenced memories rather than the neutral word pairs used in typical TNT studies. As some have correctly noted (e.g., Kihlstrom, 2002; Schacter, 2001), if the TNT paradigm is to be taken seriously as a method for studying motivated thought suppression, it must be shown to produce similar forgetting when the to-be-avoided memories are negatively valenced and more naturalistic. Perhaps, the suppression mechanisms described here are not suited for suppressing highly arousing materials. If so, this would suggest the inadequacy of this paradigm as a model for understanding suppression of negatively charged memories. In several recent studies, however, it has been shown that suppression of negative memories leads to comparable (actually non-significantly increased) inhibition relative to either neutral stimuli (Anderson & Kuhl, in preparation; Depue et al., 2006, 2007; Hart, 2006) or positive stimuli (Joorman et al., 2005; Hart, 2006; although, Hertel & Gerstle, 2003 found that non-depressed individuals were not able to suppress negative adjective–noun pairs, despite their ability to suppress positive ones). Taken as a whole, these studies suggest that negative materials can be suppressed, and may actually be suppressed more than other materials.

Another potential limitation of early TNT studies was that they relied almost exclusively on word pair stimuli. If this paradigm is meant to capture the repression of naturalistic traumatic events, then it is necessary to show that it can occur for more ecologically valid stimuli. Several studies have now shown that the observed impairments generalize to visual objects and faces (Hart, 2006) and even when the to-be-suppressed stimuli are aversive images (e.g., car accidents; Depue et al., 2006, 2007). In one experiment, Depue et al. (2006) taught subjects face-picture pairs, where the face acted as a retrieval cue to the previously studied photograph, and they manipulated whether the photograph was of a neutral scene (e.g., a horse standing in a field) or a negative scene (e.g., a car accident). After performing up to 10 "no-think" trials on these photographs, subjects were tested for all of the previously studied pairs by showing them the faces again. Subjects were asked to write a brief description of the photograph they had studied earlier in association with each face. Subjects recalled fewer of the photographs they had suppressed than baseline photographs that had not appeared in the TNT phase, mirroring the findings from studies with word pairs. As noted in the previous section, this was even true when the photographs were aversive. Thus, the suppression mechanism isolated by the TNT paradigm can be effective for suppressing complex, emotional stimuli, strengthening the claim that it is a viable model for understanding motivated memory suppression.

Some investigators have also argued that the TNT paradigm is not relevant for understanding Freudian repression because it cannot explain how individuals can have a deficit in explicit recall for a traumatic event, while still manifesting implicit memory for it (e.g., Kihlstrom, 2006). This issue has just begun to be explored, but Kawaguchi, Hotta, and Takei (2006) have now provided evidence that inhibition in the TNT paradigm can influence explicit memory, while leaving implicit memory intact. Their results demonstrate a persisting influence of inhibited information outside of awareness (see also Bjork & Bjork, 2003). This avenue of research needs further exploration, but it seems that even in this regard the TNT paradigm has promise for accounting for important aspects of repression. The research described here, however, provides no evidence to support the existence of unconscious repression, the idea that traumatic memories can be instantaneously and automatically thrust into the unconscious. Instead, we suspect that most cases of motivated forgetting arise from persisting attempts to control awareness strategically as outlined here.

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 in which people wish to avoid persistent, intrusive thoughts. By this view, repeatedly avoiding memories in naturalistic settings may cause long-lasting impairments at recalling those avoided memories. At a minimum, the recurring engagement of memory suppression in the face of intrusive remindings may be a central factor accounting for the remission of intrusiveness over time. We have argued that these findings provide an existence proof of mechanisms that could underlie repression (Anderson & Green, 2001; Anderson et al., 2004, 2006).

The body of research presented here does lie at odds with the now extensive literature on the "White Bear" phenomenon studied by Wegner and colleagues (for a review, see Wenzlaff & Wegner, 2000). In those studies, when subjects are asked to willfully prevent an unwanted thought from entering awareness (e.g., do not think of white bears), they find that the avoided thoughts are "ironically" more accessible later than non-avoided thoughts; a finding that is in direct contrast with the evidence reported here. This paradigm has proven to be quite profitable in clinical research and has been related to obsessive–compulsive disorder, depression, and anxiety. Many of these studies have found that clinical patients who show a strong tendency towards self-directed thought suppression also show the greatest tendency to suffer from intrusive thoughts. The general conclusion drawn from this body of evidence is that attempts to prevent unwanted thoughts from coming to mind are ultimately counter-productive and lead to increases in the avoided thoughts. Based on this conclusion, many clinicians have abandoned strategies involving thought suppression in favor of other methods of dealing with intrusive memories, including, in some cases, an emphasis on explicitly focusing on unwanted thoughts.

As mentioned earlier, the preceding evidence and conclusions contrast with the evidence presented in this article, namely that thought suppression can be quite effective. Why is it that subjects frequently fail at thought suppression in the white bear task, but can be successful at suppression in the TNT task? First, the difference could lie in the way the goal of the task is structured. In the white bear paradigm, the only way a subject knows whether they are "on task" or not is to think about the avoided thought ("am I thinking about white bears right now?"), whereas the same is not true for "no think" trials ("am I thinking about the word that goes with 'ordeal'?" does not require you to think of "roach"). Another potential difference is that the TNT task requires subjects to suppress a thought when presented with a very specific reminder of that thought, whereas the white bear task does not. Perhaps, the difference has something to do with a distinction between suppressing a single thought or many thoughts. These represent merely a subset of the differences between the two tasks and at present; we do not know which produces the diverging outcomes. What we suggest is that both paradigms capture actual situations where one might attempt to use thought suppression in naturalistic settings. Importantly, future work focusing on understanding the differences between these situations will help provide an understanding of when thought suppression is likely to fail and when it is likely to succeed.

Lastly, we would also like to point out that the correlation mentioned earlier between the use of thought suppression strategies and difficulty with intrusive memories does not imply anything about causality. While that correlation has been interpreted as evidence that thought suppression attempts *cause* intrusive memories, it seems just as plausible that patients with the most persistent, intrusive thoughts (due to greater symptom severity) would be the ones most likely to (perhaps desperately) engage thought suppression strategies. Given the differences in the situations captured by these paradigms and the disparate outcomes, it is important to reconsider some of the conclusions that have arisen from the research from the white bear paradigm. Although we agree that there are situations where individuals are ineffective at stopping unwanted thoughts, we disagree that such failures imply a fundamental inability to inhibit unwanted thoughts. Our paradigm indicates that people are *able* to do this. The extensive literature within cognitive neuroscience on executive control abilities suggests that such motivated thought control is not only common, but is necessary for goal-directed, intelligent thought. We anticipate that work reconciling these literatures will find that both paradigms capture aspects of naturalistic thought suppression. An understanding of why they differ will provide an invaluable contribution to the treatment of individuals who suffer from intrusive, unwanted thoughts.

1.2. Neural substrates of stopping retrieval

In addition to studying the behavioral consequences of suppressing unwanted thoughts, we have also attempted to build a neurocognitive model of how memory suppression is achieved. This model begins with the observation described earlier that stopping retrieval is analogous to stopping a motor response. We propose that stopping retrieval is accomplished by the same brain mechanisms that are recruited to stop a motor response from occurring. In the latter situation, the target of suppression is activation in motor cortical regions, while in the former situation, the target is a declarative memory representation. The actual stopping process, however, may be the same in both situations.

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 control have shown that response override is associated with a network of control-related regions, including the lateral prefrontal cortex, anterior cingulate cortex, lateral premotor cortex, and intraparietal sulcus (e.g., Garavan, Ross, Murphy, Roche, & Stein, 2002; Menon, Adleman, White, Glover, & Reiss, 2001). The lateral prefrontal regions, in particular, seem to play a critical role in stopping overt behavior. Animal research with the "go/no-go" task suggests that lesions to the homologous region in monkeys impair the ability to stop a response (Iversen & Mishkin, 1970). Even more convincing, stimulation of this region during a "go" response actually disrupts the ability to make the response (Sasaki et al., 1989). Thus, the lateral PFC seems to play a critical role in the stopping of motor responses.

If our hypothesis is correct, suppressing unwanted memories might also engage these same regions of lateral prefrontal cortex. In addition, there should also be some regions unique to memory control: those regions that are the target of suppression. Given that the goal of the TNT task is to suppress conscious recollection, a process ascribed to the hippocampus (e.g., Squire, 1992; Eldridge, Knowlton, Furmanski, Bookheimer, & Engel, 2000), this region seems a likely candidate target. Based on the foregoing analysis, Anderson et al. (2004) hypothesized that control-related regions (particularly lateral prefrontal regions, including both dorsolateral and ventrolateral 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 using a TNT task similar to the one described earlier. Replicating the earlier behavioral work, they found that subjects recalled significantly fewer suppression words than baseline words (i.e., the *negative control effect*). To investigate the hypotheses concerning neural regions underlying suppression, they compared the hemodynamic response on ‘‘think’’ trials and ‘‘no-think’’ trials. As predicted, ‘‘no-think’’ trials were associated with more activity than ‘‘think’’ trials in control-related regions that overlapped strongly with the regions involved in stopping motor responses. These results demonstrate that stopping retrieval is not simply a failure to engage retrieval processes. Instead, the fact that these control-related regions showed more activation during suppression trials than during ‘‘think’’ trials indicates that subjects must actively engage processes to prevent the unwanted memories from coming to mind. Furthermore, these findings confirm the idea that common brain regions control stopping both unwanted memories and unwanted actions.

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. This difference in hippocampal activity between ‘‘think’’ and ‘‘no-think’’ trials suggests that subjects are able to physically regulate the activity of the hippocampus to engage or disengage the recollective process, as necessitated by the current goals of the rememberer. While this difference could be explained by increased hippocampal activity during ‘‘think’’ trials, it is also consistent with the hippocampus being down-regulated during suppression. 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 below-baseline behavioral suppression (i.e., the negative control effect) is inconsistent with the idea that the difference between hippocampal activity on ‘‘think’’ and ‘‘no-think’’ trials is *entirely* due to heightened activation in the ‘‘think’’ condition. Instead, it suggests that subjects can strategically down-regulate the hippocampus to prevent conscious recollection.

This pattern of results has now been replicated by Depue et al. (2007) using the face-photograph procedure described earlier. In addition to replicating the activation of the control-related network and hippocampal down-regulation reported by Anderson et al. (2004), they also reported down-regulation of amygdala activation, suggesting an ability to control emotional responses. Critically, they also included a passive fixation condition in their experiment that allowed them to assess the normal level of activity within the hippocampus. They found that during suppression trials, the hippocampus and amygdala were not simply less engaged than during ‘‘think’’ trials, they were also significantly less active than when subjects simply stared passively at a fixation cross. This provides even stronger evidence for the idea that subjects are actually achieving memory control by down-regulating activity within the hippocampus. The results of these studies suggest a neurocognitive model of how unwanted thoughts can be regulated by executive control. Specifically, subjects prevent unwanted memories from coming to mind by the same executive control regions that are recruited to stop motor actions. Instead of targeting a motor response, these regions are recruited to suppress declarative memory representations.

1.3. Individual differences in memory control and the executive deficit hypothesis

The foregoing sections demonstrated that the ability to suppress an unwanted memory is robust and, in particular, shows a very large overall memory control effect (‘‘think’’ > ‘‘no-think’’). The meta-analysis also showed that the average magnitude of the *negative control effect* (i.e., below-baseline suppression) was a relatively modest 6%. While not large, it is also clear that there is tremendous variability in the *negative control effect*. Indeed, some individuals show *negative control effects* as large as 60%, even though the total time spent suppressing is just over one minute.[4] On the other hand, there are individuals who show a very poor ability to withhold these unwanted thoughts and actually demonstrate facilitation on the final test – in extreme cases, as large as 40%. Clearly, some people excel at this task and some people not only fail to get it right, but are also actually quite bad at suppressing thoughts. While such variability could be due to random variation or to measurement error,[5] in the next section we will consider other possible sources of this variability.

[4] Some critical commentary on the ecological validity of the TNT method for studying voluntary repression has focused on the fact that the typical negative control effect is on the order of 6–8%. However, subjects are exposed to these cues at most 16 times during our experiments for, at most, 4 s on each presentation. This means that subjects spend just over a minute (64 s) suppressing the unwanted memory. The fact that they show an 8% memory impairment after a minute of practice actually suggests a very powerful mechanism. In situations with actual traumatic memories, individuals are likely to spend much more time avoiding these memories.

[5] The most obvious way of addressing concerns about measurement and reliability is to obtain measures of test–retest reliability. This is problematic in the TNT paradigm, though, because of the surprise final memory test used to measure inhibition. If subjects were to be tested a second time, it would be difficult to ensure that they would approach the task in the same way.

If we can find stable population differences that sensibly relate to the ability to suppress, then we can reasonably assume that real individual differences do exist.

Before proceeding, we would like to point out that not every participant must show below-baseline suppression for this phenomenon to be germane to repression. Critical commentary on this research by skeptics of memory repression has focused on the "small" (Hayne, Garry, & Loftus, 2006, p. 521) nature of the negative control effect. For this paradigm to have clinical relevance, however, it is not required that *everyone* is able to suppress. All that is necessary is for there to be *some* individuals who can wield it very effectively. In the next section, we lay out a theoretical account of how variation in executive control ability might influence success at the suppression task. Importantly, this account generates abundant predictions about which factors should be related to suppression. Some of these predictions have already been confirmed, lending confidence that individual differences in memory suppression originate from variation in executive function.

1.3.1. The executive deficit hypothesis

According to the inhibitory control perspective, unwanted memories are suppressed by engaging executive control mechanisms to prevent the unwanted trace from coming to mind. If so, individual differences in executive control ability should predict variability in memory suppression. The neuroimaging work presented earlier indicates that a network of executive control regions supports memory suppression (Anderson, Ochsner, et al., 2004), and that individual differences in the recruitment of lateral prefrontal cortex predict the amount of memory inhibition. Indeed, as shown in Fig. 4, a regression analysis showed that recruitment of DLPFC strongly predicts inhibition. Given the role of executive control in achieving successful suppression, might the variability in our measure of memory suppression derive from variability in executive control more broadly? If so, then variables that influence the ability to engage executive control should predict how effectively people can suppress memories.

One task that has been used extensively to study individual differences in executive control is complex working memory span. In complex span tasks, subjects must maintain items in working memory despite engagement in a secondary task, an ability that is believed to be dependent upon executive control (e.g., Kane & Engle, 2002). If complex working memory capacity provides a good measure of executive control, then perhaps variations in working memory span would relate to the effectiveness of memory suppression. This prediction was confirmed by two recent experiments (Bell & Anderson, in preparation) showing that subjects with higher working memory capacity (defined by either verbal or visual working memory span) showed larger inhibition effects than low span subjects. Consistent with our earlier observations about the wide range of *negative control* scores, individuals with low WM capacity actually showed facilitation of the to-be-suppressed memories, whereas higher capacity individuals displayed good control over the unwanted memories (average negative control effect = 12%).

People might also differ in their ability to suppress unwanted memories because of past experience engaging in memory suppression. As we speculated earlier, exposure to trauma should put individuals in situations where they would be motivated to repeatedly suppress unpleasant remindings. If executive control abilities can be enhanced with practice, and if memory suppression in the TNT paradigm is dependent on executive control, then practice at overriding retrieval might produce greater ability to suppress unwanted memories. In two studies, Anderson and Kuhl (in preparation) found evidence consistent with this idea. People who had experienced more traumatic events (as measured by the Brief Betrayal Trauma Survey; Goldberg & Freyd, 2006) showed enhanced memory inhibition abilities when compared to individuals who had experienced little or no trauma. This establishes an important connection between inhibition in the TNT procedure, and processes likely to be used to control unwanted remindings in daily life. Furthermore, it suggests that this ability is *experience-dependent*: practice at performing memory suppression

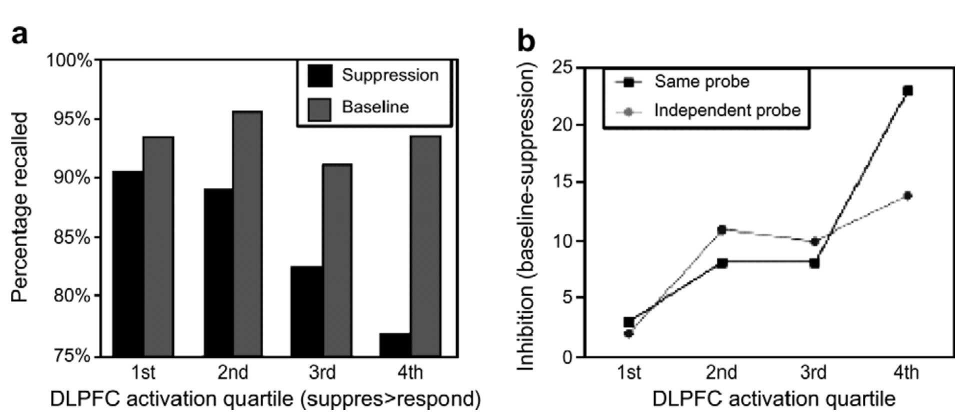

Fig. 4. Successful recruitment of DLPFC predicts behavioral inhibition. (a) Memory inhibition effects for four subject groups, differing in DLPFC activation. Subjects with greater DLPFC activity (on the right side) show reduced recall of Suppression items, but do not differ from other subjects on their recall of Baseline items. (b) Magnitude of the *negative control effect* on both the same probe and independent probe tests for each DLPFC group.

improves the ability to recruit these mechanisms even with new material unrelated to the original trauma.

According to the *executive deficit hypothesis*, in addition to variability within the undergraduate research population we should see variability in suppression across populations that differ in executive control abilities. Several prominent theories of cognitive aging claim that the primary change in cognitive abilities during normal aging is due to deteriorations in executive control abilities, particularly the ability to inhibit interference from distracting or unwanted thoughts (e.g., Hasher & Zacks, 1988). Again, the *executive control hypothesis* predicts that older and younger adults should differ in their ability to recruit executive control, with younger adults being more able to prevent unwanted memories from coming to mind. This prediction was confirmed in a recent study as a group of older adults (between the ages of 65 and 80) showed a significantly reduced ability to inhibit compared to younger adults (aged 18–25) (Anderson et al., in preparation).

People may also vary in their ability to suppress unwanted memories due to the strategies they employ. Independent of their underlying executive control abilities, the tendency for subjects to choose different strategies might influence their success at the task, especially if these strategies vary in effectiveness at inducing suppression. Most studies from our laboratory have not explicitly instructed subjects to use particular strategies; rather, we have told subjects to employ whatever strategy works for them. We have, however, collected extensive post-experimental questionnaire data about what strategies subjects report using (see Fig. 5). While a few strategies are used more often than others (e.g., generating a mental image or thinking of an alternative word), it is important to note that the *most* frequently used strategy was only used by 23% of the time, despite the fact that subjects are allowed to report using multiple strategies. Thus, there is tremendous variability in the way people approach this task. We are currently examining whether these self-reported strategies differ in terms of the amount of negative control that subjects display.

Other researchers have attempted to investigate experimentally whether certain strategies are more successful than others. Hertel and Calcaterra (2005) showed a group of "aided" subjects alternative response words for each suppression pair (e.g., for the pair "security-officer" they were shown "security-vehicle"). The "aided" subjects were then told that they could think of these alternative responses during the "no-think" trials in the TNT phase (this is similar to the "alternative word" strategy shown in Fig. 5). While the "aided" subjects showed a significant *negative control effect*, the "unaided" subjects, who were not provided with alternative thoughts, failed to show below-baseline suppression.[6] Furthermore, the "unaided" subjects who were able to suppress successfully were also more likely than the poorer inhibitors to report spontaneously generating alternative thoughts. From this pattern, Hertel and Calcaterra concluded that suppressing unwanted memories depends on the generation of alternative thoughts. While there is a certain intuitive appeal to this conclusion, it may be premature. For example, Bergström et al. (2006) instructed half of their subjects to use a "thought substitution" strategy similar to what Hertel used (they did not provide specific alternative thoughts to use, however). The other half of the subjects were given "thought suppression" instructions where they were simply told to prevent the words from coming to mind, while explicitly instructed not to generate alternative thoughts. In that study, both groups showed a negative control effect when tested using the same cue used to encode the item, unlike in Hertel and Calcaterra's study. When the suppressed memory was tested with a novel, independent cue, only the thought suppression group showed inhibition. Thus, thought suppression can not only be effective, but also appears to result in a more generalized forgetting than thought substitution.

It is clear that the majority of subjects show some degree of below-baseline suppression, despite no one strategy dominating their subjective reports (as shown in Fig. 5). This suggests that, contrary to Hertel & Calcaterra's claim, there may be many ways to achieve successful suppression. In fairness, many of the strategies shown in Fig. 5 could be classified under a broad category of "thinking of an alternative thought." In addition to thinking of alternative words, subjects report thinking of personal memories, songs, or sounds. There may be similarities between these strategies, but it is not yet clear that they all work the same way. For example, is it critical that the diversionary thought is related to the cue word in some way? What if the subject simply thought of the same diversionary thought for every suppression hint word? Would this produce comparable inhibition to thinking of a separate alternative thought for each cue word? Importantly, Hertel and Calcaterra's method of assessing spontaneously generated "alternative thought" strategies was very broad. They simply had subjects state their agreement with the statement "I kept myself from thinking about the original response word by thinking about something else." This question encompasses many of the strategies listed in Fig. 5, and ignores the considerable differences between those various strategies. Furthermore, in our analysis even the subjects who report using a "mind blanking" strategy, which would seem to be a very different strategy than thinking of a "diversionary thought," show a *negative control effect* as a group. Thus, at present, it seems unclear whether one strategy is more successful than others.

Perhaps stronger evidence against a strict "diversionary thoughts" account of successful TNT suppression comes from the neuroimaging studies described earlier. If the only thing subjects do during a suppression trial is think of alternative thoughts, we would expect activation during "think"

[6] At least some of this failure to see suppression in the "unaided" condition resulted from subjects failing to comply with the instructions. Performing a median split of the "unaided" subjects based upon their self-reported compliance yielded a sizable (yet non-significant) below-baseline suppression effect for the compliant subjects.

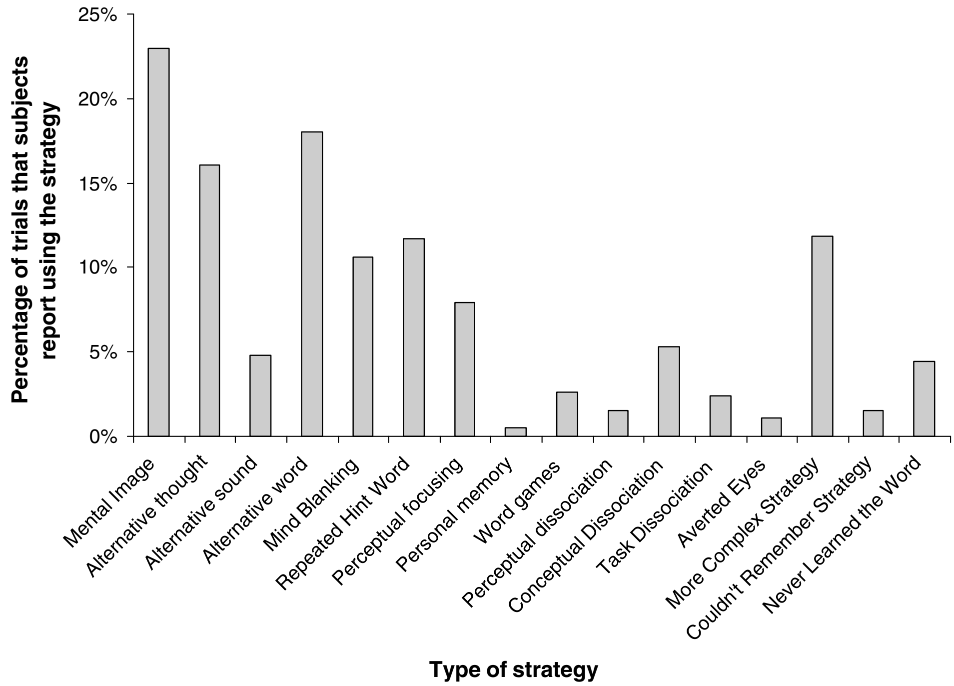

Fig. 5. Distribution of strategies employed by subjects on "no think" trials in the TNT procedure. Each column represents a different possible strategy that subjects reported using on post-experimental questionnaires. Subjects were allowed for each "no think" hint word to select as many options as fit their strategy for that particular word. The *y*-axis shows the percentage of trials, averaged across subjects and items, that each strategy was employed.

and "no-think" trials, to look quite similar, as both tasks involve retrieving an item from long-term memory. However, both Anderson et al. (2004) and Depue et al. (2007) found significantly less hippocampal activation on "no-think" trials than "think" trials. If subjects are simply thinking of alternative thoughts on "no-think" trials then they should be engaging the hippocampus to the same degree on "think" and "no-think" trials. Supporting this notion, recent neuroimaging evidence using the retrieval practice paradigm shows that inhibiting thoughts *without* a conscious intention to suppress does not yield a decrease in hippocampal activity (Kuhl, Dudukovic, Kahn, & Wagner, 2007). If retrieval-induced forgetting fails to down-regulate hippocampal activity, then it fails as a candidate mechanism for explaining the hippocampal down-regulation observed in TNT studies. To be clear, we are not suggesting that retrieval of diversionary thoughts (i.e., retrieval-induced forgetting) plays no role in producing memory impairments in the TNT task. Rather, we suggest that it cannot account for the observation of hippocampal down-regulation in TNT studies. More generally, we suggest that suppression of unwanted memories is not achieved through simply retrieving alternative thoughts.

The preceding examples represent only a subset of the predictions that arise from the *executive deficit hypothesis*. For example, recent research has suggested that there are genetic contributions to executive control abilities (e.g., Fossella et al., 2002). Perhaps, genetic variation in the production of neurotransmitters related to executive-control may help explain variability in the *negative control effect*. One would also expect differences in executive control ability across different physiological and psychological states. For example, sleep deprivation is known to impair the ability to perform tasks that depend on executive control (e.g., Pilcher & Huffcutt, 1996). Chronic sleep deprivation might be an especially important factor underlying variability in performance, especially since the undergraduate students who normally participate in these studies are notoriously variable in the amount of sleep they get. Because traumatic experiences often disturb normal sleep patterns, sleep deprivation may be a contributing cause to the early prevalence of intrusive memories in the aftermath of trauma. The *executive deficit hypothesis* also predicts that suppression ability should be diminished whenever people are in experimental conditions that tax their executive control abilities (e.g., under divided attention conditions). These examples illustrate the diversity of questions generated by the executive deficit hypothesis, and suggest that it provides a profitable framework for guiding research on the management and recovery from intrusive memories in the aftermath of trauma.

1.3.2. Conclusions and future directions

After experiencing a traumatic event, people are often troubled by intrusive memories. Some people succeed in putting these unwanted memories out of mind, whereas others seem to be unable to do so. In this paper, we have argued that the capacity to control unwanted remindings in the face of unwelcome reminders arises from general executive control processes of the sort that are widely studied in cognitive neuroscience today. In support of this, we reviewed evidence from our laboratory model of memory suppression and have provided examples of how individuals differ in their ability to perform this task. Based on these findings, we have argued that individual differences in the ability to cope with unwanted memories can, to a large degree, be explained by differences in executive control. Although other factors may contribute to memory control, this framework provides a rich model that helps to understand this function at the cognitive and neurocognitive level. We believe that over time, this model will help us to understand the differing experiences of people coping in the face of trauma – both people suffering from intrusive memories (e.g., PTSD patients) and people who have genuine memory lapses in the ability to recall unwanted memories.

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. (2006). Repression: A cognitive neuroscience approach. In M. Mancia (Ed.), *Neuroscience and Psychoanalysis* (pp. 327–350). Milan: Springer.

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 & T. H. Carr (Eds.), *Inhibitory processes in attention, memory and language* (pp. 265–326). San Diego, CA: 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*, 366–369.

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

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

Anderson, M. C., & Levy, B. J. (2002). Repression can (and should) be studied empirically: Reply from Anderson and Levy. *Trends in Cognitive Sciences, 6*, 502–503.

Anderson, M. C., & Levy, B. J. (2006). Encouraging the nascent cognitive neuroscience of repression. *Behavioral and Brain Sciences, 29*, 511–513.

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.

Aslan, A., Bäuml, K., & Pastötter, B. (2007). No inhibitory deficit in older adults' episodic memory. *Psychological Science, 18*, 72–78.

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.

Bäuml, K., & Hartinger, A. (2002). On the role of item similarity in retrieval-induced forgetting. *Memory, 10*, 215–224.

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.

Bergström, Z., Richardson-Klavehn, A., & de Fockert, J. (2006). Electrophysiological correlates of thought-substitution and thought-suppression strategies in the think/no-think task. Paper presented at the 4th international conference on Memory, Sydney.

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

Bulevich, J. B., Roediger, H. L., Balota, D. A., & Butler, A. C. (2006). Failures to find suppression of episodic memories in the think/no-think paradigm. *Memory and Cognition, 34*, 1569–1577.

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

Carroll, M., Campbell-Ratcliffe, J., Murnane, H., & Perfect, T. (2007). Retrieval-induced forgetting in educational contexts: Monitoring, expertise, text integration, and test format. *European Journal of Cognitive Psychology, 19*, 580–606.

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

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.

de-Zubicaray, G. I., Andrew, C., Zelaya, F. O., Williams, S. C. R., & Dumanoir, C. (2000). Motor response suppression and the prepotent tendency to respond: A parametric fMRI study. *Neuropsychologia, 38*, 1280–1291.

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. (1915/1963). Repression, trans. C. M. Baines & J. Strachey. In: Strachey, J., (Ed.), *The standard edition of the complete psychological works of Sigmund Freud* (*Vol. 14*, pp. 146–58). London: Hogarth Press.

Fossella, J., Sommer, T., Fan, J., Wu, Y., Swanson, J. M., Pfaff, D. W., & Posner, M. I. (2002). Assessing the molecular genetics of attention networks. *BMC Neuroscience, 3*, 14.

Garavan, H., Ross, T. J., & Stein, E. A. (1999). Right hemispheric dominance of inhibitory control: An event-related functional MRI study. *Proceedings of the National Academy of Science, 96*, 8301–8306.

Garavan, H., Ross, T. J., Murphy, K., Roche, R. A. P., & Stein, E. A. (2002). Dissociable executive functions in the dynamic control of behavior: Inhibition, error detection, and correction. *Neuroimage, 17*, 1820–1830.

Goldberg, L. R., & Freyd, J. J. (2006). Self-reports of potentially traumatic experiences in an adult community sample: Gender differences and test–retest stabilities of the items in a brief betrayal-trauma survey. *Journal of Trauma and Dissociation, 7*, 39–63.

Hart, R. E. (2006). Applications of intentional suppression: The effect of suppression on recognition memory of ecologically valid stimuli. Unpublished doctoral dissertation.

Hasher, L., & Zacks, R. T. (1988). Working memory, comprehension, and aging: a review and a new view. *The psychology of learning and motivation, 22*, 193–225.

Hayne, H., Garry, M., & Loftus, E. F. (2006). On the continuing lack of scientific evidence for repression. *Behavioral and Brain Sciences, 29*, 521–522.

Hertel, P. T., & Calcaterra, G. (2005). Intentional forgetting benefits from thought substitution. *Psychonomic Bulletin and 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 and 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.

Iversen, S. D., & Mishkin, M. (1970). Perseverative interference in monkeys following selective lesions of the inferior prefrontal convexivity. *Experimental Brain Research, 11*, 376–386.

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.

Kane, M. J., & Engle, R. W. (2002). The role of prefrontal cortex in working-memory capacity, executive attention, and general fluid intelligence: An individual-differences perspective. *Psychonomic Bulletin and Review, 9*, 637–671.

Kawaguchi, J., Hotta, C., & Takei, S. (2006). Implicit memory for intentionally suppressed information. Poster presented at the 18th annual meeting of the association for psychological science, New York, NY, May 25–28.

Kihlstrom, J. F. (2002). No need for repression. *Trends in Cognitive Sciences, 6*, 502.

Kihlstrom, J. F. (2006). Repression: A unified theory of a will-o'-the-wisp. *Behavioral and Brain Sciences, 29*, 523.

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. (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. (2007). Inhibiting your native language: The role of retrieval-induced forgetting during second language acquisition. *Psychological Science, 18*, 29–34.

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

Menon, V., Adleman, N. E., White, C. D., Glover, G. H., & Reiss, A. L. (2001). Error-related brain activation during a go/no go response inhibition task. *Human Brain Mapping, 12*, 131–143.

Pilcher, J. J., & Huffcutt, A. I. (1996). Effects of sleep deprivation on performance: A meta-analysis. *Sleep, 19*(4), 318–326.

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

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

Sakagami, M., & Niki, H. (1994). Spatial selectivity of go/no-go neurons in the monkey prefrontal cortex. *Experimental Brain Research, 100*, 165–169.

Sasaki, K., Gemba, H., & Tsujimoto, T. (1989). Suppression of visually initiated hand movement by stimulation of the prefrontal cortex in the monkey. *Brain Research, 495*, 100–107.

Saunders, J., & MacLeod, M. D. (2006). Can inhibition resolve retrieval competition through the control of spreading activation? *Memory and Cognition, 34*, 307–322.

Schacter, D. L. (2001). Suppression of unwanted memories: Repression revisited? *The Lancet, 357*, 1724–1725.

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.

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 and Cognition, 32*, 602–609.

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.

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

Verde, M. F. (2004). The retrieval practice effect in associative recognition. *Memory and Cognition, 32*, 1265–1272.

Wenzlaff, R. M., & Wegner, D. M. (2000). Thought suppression. *Annual Review of Psychology, 51*, 59–91.

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.