

Inhibition in memory

Jeroen G.W. Raaijmakers

UNIVERSITY OF AMSTERDAM

Abstract

In the past 20 years, a new approach to forgetting has been proposed, based on the notion of inhibition. According to this view forgetting is partly due to a process of inhibitory control. In this chapter, I review the current status of this account, discussing the evidence that has been proposed as well as the counterarguments that have been made. I show that the paradigms that are likely to generate the best evidence for inhibition are the retrieval induced forgetting paradigm and the think / no-think paradigm. Other paradigms such as directed forgetting and part-list cuing seem to be better explained using other memory principles. I conclude that there is a need for a consistent and preferably formalized model that combines inhibition and competition as principal factors in forgetting.

Key Terms: Inhibition – Forgetting – Interference – Retrieval Induced Forgetting – Inhibitory control – Repression – Memory models.

Introduction

Clive Wearing was a highly successful British musician and conductor who at the height of his career, at an age of 46, contracted herpes encephalitis. While he recovered from the disease itself, his brain and in particular the hippocampal region was severely damaged. As a result of this damage, he developed a dense amnesia, being unable to recall events that happened just 20 seconds before. Clive's case has featured in several documentaries and, if you ever watch one, you will agree that his current life is not a happy one. Indeed, most people believe that a good memory is desirable. However, what if your memory is so good that you are basically unable to forget past events? Although this may sound like science fiction, there are in fact a number of documented cases with such an exceptional memory. The first of these was discovered when a Californian woman, Jill Price, contacted James McGaugh asking for help. She described her problem as having no control over her memory: "I run my entire life through my head every day and it drives me crazy!!!" and "It is non-stop, uncontrollable and totally

exhausting...” (Parker, Cahill, & McGaugh, 2006, p. 35). Clearly, this is not a very enviable condition.

These two extremes illustrate that for a ‘good’ memory it is not enough to be able to store information in memory: We also need some control over what is retrieved at any specific moment. According to standard memory theories, such control is at least partly achieved through the use of retrieval cues, other information that helps to focus the search process and that (loosely speaking) filters out irrelevant information to be able to retrieve the target information. For example, in the SAM theory for memory retrieval (Raaijmakers & Shiffrin, 1981) it is assumed that adding an extra retrieval cue will increase the likelihood of retrieving memory traces that are associated to that cue and decrease the likelihood of retrieving memory traces that are not associated to that cue. In SAM, the associative strengths to the individual cues are multiplied to compute an overall associative strength for the combined probe set. In such a model, if a specific trace is completely unassociated (i.e., associative strength is zero), that trace will have a probability of zero of being sampled using those retrieval cues, and hence effectively drops out of the set of retrievable items.

The basic question that the present chapter deals with is whether such a mechanism is the only mechanism available or whether there are alternative mechanisms by which memory control is achieved. To make things more concrete, suppose a cue Q_1 is associated to a number of memory traces $A_1, A_2, A_3, A_4,$ and A_5 . Another cue, Q_2 , is associated to some of the same memory traces and to some others: $A_2, A_5, A_6, A_7,$ and A_8 . According to the rule described above, the combined cue, $Q_1 \oplus Q_2$, will be associated to A_2 and A_5 . By using the cue Q_2 in addition to Q_1 we have narrowed down the search set to these two memory traces. In such models, it is usually assumed that the likelihood of retrieving each of these two items on a recall test is determined by their relative strengths to the combined probe cue. For example, if the strength of A_2 is 1.5 (arbitrary units) and the strength of A_5 is 3.0, then the likelihood of retrieving A_2 is 0.33. Hence, the likelihood of retrieval of A_2 is reduced due to the competition from A_5 . According to this *competitive retrieval assumption*, forgetting is at least partly due to the competition from other memory traces associated to the set of probe cues.

Now suppose that the association of A_5 to the cue set is quite a bit stronger than that of A_2 but that it is A_2 that we would like to retrieve. Is there some way in which the system might be able to make the retrieval of A_2 more likely, other than simply repeated retrieval attempts (repeated sampling)?

According to one group of researchers, the answer to that question is ‘yes’ and the mechanism that they have proposed is ‘*inhibition*’ or ‘*suppression*’. The basic idea is that the trace of A_5 that is blocking or interfering with the retrieval of A_2 is suppressed or reduced in strength, thereby making the retrieval of A_2 more likely. This inhibition proposal is similar in some respects to the older Freudian

notions of suppression. The basic difference is that according to Freudian notions the suppression is due to an emotional need to keep the information out of consciousness whereas in these modern proposals the suppression is not due to an emotional need but to the actual act of trying to forget (inhibit) that information. That is, the assumption is that during the attempt to retrieve A2, the (stronger) item A5 might be activated or retrieved and that subsequently A5 is suppressed (since it is recognized to be incorrect) and this suppression of A5 makes it less likely to be retrieved on the next retrieval attempt. In this way, the likelihood of retrieving the target item A2 is increased. A crucial assumption of this inhibition theory is that the suppression of A5 is not just a temporary effect but one that endures for some time (how long exactly is not clear but it should last for at least several minutes or even hours).

This idea, that the strength of a competing trace is reduced while trying to retrieve the target trace, has a long history that dates back to at least the 1940s when Melton and Irwin (1940) proposed that competition alone was not enough to explain the results of their experiments on retroactive interference and that an additional factor had to be involved (which they initially called Factor X). This additional factor was soon termed *unlearning*, by analogy to the unlearning that was known in the animal conditioning literature. The assumption was that during the learning of a new association A-C a previously learned association A-B would be weakened or unlearned. The most important evidence for this assumption was that retroactive interference was still obtained on a so-called MMFR test where the participants were allowed to give both the B and the C response (presumably eliminating the effect of the competition between these two responses). Although this assumption was later shown to be incorrect (see Mensink & Raaijmakers, 1988), the idea that interference and forgetting could be explained by such a Two-Factor Theory based on competition and unlearning remained popular for many years. Despite its popularity, the notion of unlearning was abandoned around 1975 when the problems it had in explaining the results of several experiments became insurmountable. For example, within the Two-Factor framework, proactive interference was the result of competition only, yet it could also be obtained with MMFR testing that was claimed to eliminate the effect of competition (for a review, see Crowder, 1976; Mensink & Raaijmakers, 1988).

As I will describe in more detail later in this chapter, the notion of inhibition re-emerged in later years, initially just as a descriptive term (see e.g., Roediger, 1974, 1978) but later also as a causal factor in explaining forgetting in various experimental paradigms (e.g., Bjork, 1989). More recently, several authors (Anderson, Bjork, & Bjork, 1994; Anderson, 2003; Bäuml, 2008) have proposed that inhibition should be seen as the major cause of forgetting. It is these newer proposals that will be the focus of the present chapter. In the next sections, I will first describe the standard non-inhibitory account of forgetting and the inhibition account of forgetting as proposed by Anderson (2003) and Bäuml (2008),

followed by a more detailed review of the evidence that has been put forward in recent years. Although the discussion in this chapter focuses on inhibition (versus competition) as an explanation for forgetting, it should be mentioned that these are not the only mechanisms that have been proposed to explain forgetting. For example, several researchers (see Wixted, 2004) have proposed consolidation as a factor that contributes to forgetting (in addition to competition). The basic idea is that long-term memory storage involves a process of consolidation that takes some time and that new learning may interfere with the successful consolidation of previously learned information. This hypothesis has received a lot of support (especially from neurobiological research) but an extensive discussion is beyond the scope of this chapter.

Alternative accounts of forgetting

Competition-based accounts

A standard assumption in many theories of memory (e.g., the SAM theory, Raaijmakers & Shiffrin, 1981; Mensink & Raaijmakers, 1988) is that information that was once stored in memory will remain in memory and that forgetting reflects retrieval failure. That we fail to recall specific information because it is (temporarily) inaccessible rather than unavailable is a common belief among psychologists (Loftus & Loftus, 1980) but one has to keep in mind that it is a *belief* rather than a proven fact. There is no way in which such a belief could be shown to be correct or incorrect. In the framework that we have been using, it is merely a convenient assumption that fits the observation that we often are able to retrieve a seemingly lost memory, given the right retrieval cues. The present discussion of competition-based accounts of forgetting will be based on this SAM (or SAM-REM as it is now often called) framework of memory. It should be noted, however, that the discussion is not specifically tied to this framework and that there exist similar frameworks that make more or less identical predictions for interference paradigms (e.g., John Anderson's ACT-R theory, see Anderson, 1993; Anderson, Bothell, Lebiere, & Matessa, 1998).

In such a framework, there are two basic factors that are responsible for any observed forgetting (i.e., a failure to retrieve a memory trace at time B that was retrievable at time A):

1. *A decrease in the relative strength of the memory trace.* This might be due to an increase in the number of other traces associated to the probe cues or to an

increase in the strength of these other traces, or both. This may occur even though the probe cues and hence the associative strength of the target trace to those probe cues has not changed. For example, if some of the other traces are strengthened through additional study trials, this will lead to a decrease in the relative strength of the target trace. Similarly, if some of the other items have been tested in the interval between A and B, this will also decrease the relative strength of the target trace since it is assumed that a successful retrieval will lead to a relatively large increase in the strength of the tested item (the so-called testing effect).

2. *A decrease in the associative strength of the target trace to the probe cues.* Such a decrease may occur due to a change in the probe cues from time A to time B. For example, the context may have changed such that the context at time B does not match the context stored in the memory trace as well as the context at time A. Such a context change may occur due to more or less gradual, random fluctuations in a single experimental session (as in the contextual fluctuation model, see Mensink & Raaijmakers, 1988; 1989). But context may also have changed because of external factors (such as a change in the environmental context) or because of a strategic change in the mental context by the person him/herself. For example, in an experiment in which several consecutive lists are studied and tested, the participant may create a new mental context at the start of a new list in order to minimize interference from the previously studied lists.

A key aspect of such a model is that recall is determined not just by the absolute strength of a memory trace but also by its relative strength (relative to all other memory traces). This would appear to be a well-established principle for which there is ample evidence in the literature. For example, Wixted, Ghadisha, and Vera (1997) showed that such a model accurately predicts both the probabilities and the latencies of recall for lists varying in the proportion of weak and strong items. Although in pure-strength lists, strong items are recalled as fast as weak items, in mixed-strength lists the strong items are recalled significantly faster than the weak items and faster than strong items from a pure-strength list, just as predicted from a relative-strength model. Such a pattern of results seems difficult to explain by any model in which relative strength is irrelevant to recall.

The general model as described above is designed to explain phenomena in recall paradigms, paradigms where the goal is to retrieve a specific target. In recognition paradigms, where one only has to judge whether the test item did or did not occur in a specified context, it is (generally speaking) not necessary to retrieve a specific memory trace, all that is needed is some sense of confidence or familiarity that it did occur in that context. According to most current models for recognition, such a decision may be based on the global familiarity or activation from memory. That is, if the global activation is high, we may conclude that it probably did occur in that context and if it is low we conclude that it probably did

not occur. Of course, the target item that is presented may lead to the retrieval of a specific memory trace, but the crux is that even when it does not, we may still be able to make a judgment that will often be correct.

The notion that recognition is based on global familiarity has a number of advantages. First, it provides a simple explanation for the observation that we are often able to give a quick ‘no’-response: There is no need to do an exhaustive search to find a memory trace for the test item; if the familiarity is low, a quick ‘no’-response may be given. Second, and more relevant for the present discussion, it may provide an explanation for why interference effects that are standard in recall paradigms sometimes do and sometimes do not show up in recognition paradigms. For example, although increasing the number of items on the list decreases performance (see Annis, Lenes, Westfall, Criss, & Malmberg, 2015), increasing the strength of the other items on the list does not decrease performance (the so-called list-strength effect). This somewhat paradoxical finding may be explained by assuming that the extent to which other items are activated depends not only on the number of stored features that match the presented test item but also on the number of mismatching features: The activation increases with the number of matching features but decreases with the number of mismatching features (see the REM model proposed by Shiffrin and Steyvers, 1997). Increasing the strength of other list items increases the number both of matching and of mismatching features, leading to a null effect on the observed recognition performance (see Shiffrin, Ratcliff, & Clark, 1990).

It should be noted that the null list-strength effect depends on the assumption that the manipulation that increases the strength increases the previously stored trace and does not lead to the storage of a new trace. If a new trace is formed (as might be the case if there is a change in context), this newly formed trace will increase the amount of interference. Another condition in which strengthening a competitor item will decrease recognition performance for the target item is when the strengthening has a larger effect on the matching features compared to the mismatching features, as might be the case when additional practice focuses on features that are common between targets and competitors.

Inhibition-based accounts

The starting point for the inhibition account of forgetting is the assumption that we have some control over what is being retrieved from memory. Anderson (2005) uses the analogy with response inhibition to explain the basic idea. If a plant falls off the sill while opening a window, we automatically initiate a response sequence to catch the falling plant. However, if we notice that it happens

to be a cactus plant, we may stop or override this automatic response by a process of response inhibition. This fits with the general definition proposed by MacLeod (2007, p. 5): “Cognitive inhibition is the stopping or overriding of a mental process, in whole or in part, with or without intention.”

According to Anderson (2003, 2005), a similar cognitive control process might be at work in memory retrieval when there is a strong memory trace connected to a cue that would normally be retrieved but that is not the one that we are looking for at this time. For example, one might be trying to retrieve the family name of an actor with the first name ‘George’ who played in an episode of the TV series Columbo. In that case, the cue George might lead to retrieval of the incorrect response ‘Clooney’. According to the inhibition account, we invoke a cognitive control process of inhibition to suppress the response ‘Clooney’ to be able to retrieve the correct response ‘Hamilton’. Hence, just as one is able to suppress a strong motor response, one might also be able to suppress a strong but inappropriate memory to increase the likelihood of retrieving the memory that one is looking for.

A crucial additional assumption is that this suppression of an inappropriate memory trace will have a long-term effect: The suppressed memory trace is not just momentarily weakened but the probability of activating the trace will be reduced for up to several days (Storm, Bjork & Bjork, 2012). In this respect, the memory inhibition effect is decidedly different from motor response inhibition where the effects disappear after about 20 trials (Verbruggen & Logan, 2008). Such a longer-lasting inhibition effect is of course also required if this type of suppression is to play a role in explaining phenomena such as ‘motivated forgetting’ and recovered memories (Anderson, 2001, 2006; Anderson & Huddleston, 2012; Storm et al., 2015).

A major difference between this inhibition account and the competition-based accounts is that the inhibition account assumes that the factors that are responsible for the forgetting observed on a later test operate during the previous attempts at retrieval of other, related, items. It is the retrieval or attempted retrieval of the related items that causes the forgetting. For that reason, this type of forgetting is often referred to as *retrieval-induced forgetting* (RIF). Competition-based accounts on the other hand maintain that the forgetting is largely caused by factors operating at the time of final testing.

The inhibition account as proposed by Anderson (2003) makes a number of assumptions that differentiate it from other models of memory. First, it assumes that it is not the association between the cue and the memory trace that is inhibited but that it is the memory trace itself that is inhibited. Hence, it should not matter how the item is tested, inhibition should be observed on all tests. Consistent with this assumption, retrieval induced forgetting has been observed not just in recall-based tests but also in recognition and sometimes even in implicit memory

paradigms (Veling & Knippenberg, 2004). It has also been observed in extralist cuing tasks where at the final test a new (not previously seen) word that is weakly associated to the target item, is presented as a cue (e.g., the cue ‘honey’ for the target word ‘bear’).

Second, the inhibition account tends to minimize the role of competition in tests where item-specific cues are used, such as recognition or cued recall where the (unique) initial letters of the target item are presented as cues. In its most extreme form, it is assumed that competition has no effect on the probability of recall although it may affect the latency of recall (see Shivde & Anderson, 2001, p. 176). In less extreme versions (see e.g., Storm & Levy, 2012), it is acknowledged that both inhibition and competition contribute to some degree to all demonstrations of retrieval-induced forgetting.

Paradigms used to test inhibition

The Retrieval Practice paradigm

The most frequently used experimental paradigm to investigate retrieval-induced forgetting is the Retrieval Practice paradigm, originally developed by Anderson, Bjork, and Bjork (1994). The design is illustrated in Figure 8.1. In the initial study phase, the participants are presented a list of word pairs, each pair consisting of a category name and the name of an exemplar from that category. Usually there are about 8 categories each consisting of about 6 exemplars. After one or two presentations of the study list, the second phase of the experiment is started where some of the word pairs are given retrieval practice. In this retrieval practice phase, the participant has to retrieve some of the previously presented exemplars using the category name and the initial one or two letters as cues. However, not all of the categories are presented in this retrieval practice phase, so there are practiced and nonpracticed categories (RP and NRP). Furthermore, only some of the items from the RP categories are practiced. The practiced exemplars are denoted as RP+, and the non-practiced items from these categories are denoted as RP-. Hence, in these experiments, there are three types of items: RP+, RP-, and NRP items.

According to the inhibition account, the other (unpracticed) items from the practiced category will initially be blocking or interfering with the retrieval of the




Figure 8.1: Standard design of a retrieval induced forgetting experiment. In the initial study phase, a series of category-exemplar pairs is studied. Next, some of the items from some of the categories (in this case FRUIT) are given repeated retrieval practice trials in which the category name as well as the initial two letters from the target item are presented as cues. Other categories are not practiced at all. In the final test phase, all of the category-exemplar pairs are tested by presenting the category name and the first letter as cues. In this example, FRUIT is a practiced category and *orange* and *kiwi* are practiced items (RP+ items). *Apple* is a nonpracticed item from a practiced category (a RP- item). ANIMAL is a nonpracticed category for which none of the items were practiced (NRP items).

target RP+ items. To overcome this interference, these unpracticed items will be suppressed. After a number of such practice trials, the RP+ items will have become stronger (the practice effect) and the RP- items will be weaker (the inhibition effect), all in comparison to the NRP items. Inhibition will manifest itself on a final test of all studied items, usually cued recall with the category name and the initial letter as cues.

Such a retrieval-induced forgetting (RIF) effect has been demonstrated in many experiments and is by now a well-established phenomenon. The basic finding can be readily explained by the inhibition account but there are alternative explanations that do not rely on the notion of inhibition and that are more in line with the traditional theories of forgetting. The most important ones are based on retrieval competition and context change (or a combination of these two). The competition account (see Raaijmakers & Jakab, 2013) assumes that the retrieval of the RP- items suffers from competition by the now stronger RP+ items while the NRP items do not have such strengthened competitors. The context change account (see Jonker, Seli, & MacLeod, 2013) assumes that there is a change in context between the original study phase and the retrieval practice phase such that the RP categories get associated with the retrieval practice context while the NRP categories remain connected to the original study phase. At the final test, the RP category cue reactivates the retrieval practice context, while the NRP category cues reactivate the original study context. As a result there is a mismatch for the RP- items between the context at test (the retrieval practice context) and the

context that is stored in the traces of the RP- items (the original study context). This mismatch does not occur for the NRP items and therefore recall for the RP- items is lower than that for the NRP items.

There are a number of properties, however, that have been claimed to provide unique support for the inhibition account (Anderson, 2003). These properties are interference dependence, strength independence, retrieval specificity, and cue independence. I will discuss each of these together with the relevant evidence.

Interference dependence. This property refers to the prediction that the decrease for the RP- items should be dependent on the extent to which those items are interfering during the retrieval practice of the RP+ items. Only strong competitors will need to be suppressed, and so only those items should show inhibition. Anderson, Bjork, and Bjork (1994) provided support for this prediction. In their experiments, categories were presented that consisted of either all strong or all weak exemplars. A significant RIF effect was only obtained for the strong categories. According to Anderson et al. (1994, p. 1066), a standard competition-based account would predict the opposite result: more forgetting for the weak categories. However, simulations with the SAM model showed this claim to be incorrect: such models will under most conditions predict an effect that is about equal for the strong and weak categories (see Jakab & Raaijmakers, 2009, p. 608).

Storm, Bjork, and Bjork (2007) used a variation of the retrieval practice paradigm in which, after the initial study phase, either a 'forget' or a 'remember' cue was given, indicating whether the previously studied list would later be tested or not. At the end of the experiment, an unexpected final test was given for all items (including the ones they were told to forget). A RIF effect was only obtained for the items that the participants had to remember, a result that might be explained by the assumption that these items remained strong and therefore needed to be suppressed during the intermediate retrieval practice phase. However, as I will discuss later, such directed forgetting effects are better conceptualized as being due to context change and such an explanation would fit this specific result as well.

In a cleverly designed experiment, Chan, Erdman, and Davis (2015) obtained rather strong supporting evidence for the interference dependence assumption. In their experiments, the nonpracticed (RP-) items were presented either before or only after the retrieval practice phase. Obviously, items presented only after the retrieval practice phase cannot compete during the retrieval practice and hence there should be no RIF for those items. This was exactly what they observed. Although these results provide clear prima facie evidence for the inhibition hypothesis, alternative explanations cannot be ruled out. First, it seems likely that the additional strength stored during the retrieval practice phase depends on the difficulty of the retrieval: If the competitor items have not yet been

presented, the retrieval will be easier and this may lead to less additional storage. Chan et al. (2015, p. 1309) do consider this alternative interpretation but reject it based on the complete absence of RIF in that condition. However, as we will see shortly in the section on retrieval specificity, such an argument is flawed because observed recall probability is not good indicator for the stored strengths and a substantial increase in RP+ recall can coexist with the absence of a RIF effect for the RP- items. Second, the context change hypothesis might explain the results if it is assumed that the context reactivated at test for the RP categories will be more similar to the context during the study of the RP- items when these items are presented after the retrieval practice phase because that context is now the more recent one (see Chan et al., 2015, p. 1309-10).

A number of researchers have, however, failed to obtain evidence in support of the interference dependence assumption. For example, the results of Anderson et al. (1994) could not be replicated by Williams and Zacks (2001). In their experiment the magnitude of the RIF effect was similar for the weak and strong categories, in line with the SAM prediction. Jakab and Raaijmakers (2009) obtained a similar result when strength was manipulated either by additional study trials or by selecting items based on the serial within-category position (with early items being stronger than later items).

Proponents of the inhibition account have argued that the results of Williams and Zacks (2001) and Jakab and Raaijmakers (2009) are due to the fact that these researchers used a final recall test in which only the category name was presented rather than the category name plus the initial one or two letters (Storm, 2010; Storm & Levy, 2012; Chan et al., 2015). The problem with such a category-cued recall test is that there is no control for output interference, i.e., the stronger RP+ items will probably be output first and this will exaggerate the RIF effect. According to this critique, the RIF effect for the strong items would be due to inhibition whereas the RIF effect for the weak items would be due to output interference on the final test. However, this objection does not appear to be a valid one. First, in these experiments, the same procedure was used as in Anderson et al. (1994). It is unclear why output interference would lead to equal effects in the experiments of Williams and Zacks (2001) and Jakab and Raaijmakers (2009) but not in the experiments of Anderson et al. (1994). Second, output interference should affect all items, not just the weak ones. If anything, the effect should be much larger for the strong items than the weak ones (as was also borne out in simulations with the SAM model, see Raaijmakers & Jakab, 2013) since strong RP- items should shift to a later output position much more than weak RP- items.

In addition, there are quite a few experiments that show RIF with items that should be considered weak (Perfect et al., 2004; Verde, 2012). In a recent experiment, Tempel and Frings (2015) had participants learn associations between arbitrary category names and novel pronounceable letter strings (novel 'words') and then carried out a standard retrieval induced forgetting manipulation.

Significant RIF was obtained. Although the authors concluded that this finding violated the interference dependence principle, one could also argue that whereas the semantic associations were (obviously) very weak (i.e., nonexistent), the episodic associations were probably quite strong and these may have been responsible for the observed RIF effect. Such a conclusion would seem to fit the general inhibition framework that assumes that any incorrect memory traces that intrude during retrieval will have to be suppressed, whether these are semantic or episodic.

Retrieval specificity. According to the inhibition account, RIF is due to suppression of the RP- items during the retrieval practice of the RP+ items. Such a suppression of other, related, items is assumed to be required only when there is an interference problem that has to be resolved. Hence, the RIF effect is specific to tasks in which the target RP+ item has to be retrieved from a set of competitors that includes the RP- items. Any other type of practice would not lead to a RIF effect. This is what is meant by the statement that RIF has the property of retrieval specificity.

Tests for this assumption are usually based on a comparison between two conditions: the standard retrieval practice condition and a condition in which the RP+ items are also strengthened but not in a way that involves competitive retrieval, e.g., by additional study presentations or by having the participants retrieve the category name given the exemplar as a cue (rather than the other way around). In both cases, there should be no competition from related exemplars, hence no need to resolve any interference, and therefore no suppression would be required. According to competition-based accounts, however, it should not matter how the RP+ items have become strengthened, all that matters is that they are now stronger.

A typical experiment in this line of research was carried out by Anderson, Bjork and Bjork (2000). They varied the nature of the retrieval practice task. In the competitive practice condition, the standard retrieval practice task was used. In the noncompetitive practice condition, instead of giving the category name and the first letters of the target item, they gave the target item itself and the participants had to generate the category name (e.g., FR--- orange). They obtained a pattern of results that has since then been replicated many times (see Figure 8.2a). Both conditions showed about equal levels of RP+ recall but RIF was obtained only in the condition that involved competitive retrieval, just as predicted by the inhibition account. Similar results have been obtained in other experiments in which the noncompetitive practice condition was replaced by additional study trials for the RP+ items (e.g., Anderson & Bell, 2001; Bäuml, 2002; Bäuml & Aslan, 2004; Ciranni & Shimamura, 1999; Gómez-Ariza, Fernandez, & Bajo, 2012; Hanslmayr, Staudigl, Aslan, & Bäuml, 2010; Hulbert, Shivde, & Anderson, 2012; Johansson, Aslan, Bäuml, Gäbel, & Mecklinger, 2007; Staudigl, Hanslmayr, & Bäuml, 2010;




Figure 8.2: Observed (Panel A) and predicted (Panel B) results for the experiment of Anderson, Bjork, and Bjork (2000). The predicted results in Panel B were generated from a simple non-inhibitory model based on the SAM model of Raaijmakers and Shiffrin (1981) as discussed by Raaijmakers and Jakab (2012).

Wimber, Rutschmann, Greenlee, & Bäuml, 2009). In all of these studies, RIF was obtained only with the standard retrieval practice condition.

Raaijmakers and Jakab (2012), however, provided simulation results using the SAM model to demonstrate that such a result is in fact quite compatible with a competition-based account. The crux of their analysis rests on the observation that the observed level of recall is not a good measure for the underlying strengths of the memory traces. Equal levels of recall may correspond to widely different distributions of strengths. An often overlooked aspect of the standard retrieval induced forgetting paradigm is that during the retrieval practice task no feedback is given. This implies that successive practice trials will be highly dependent. If an item has not been retrieved on the first practice trial, it is unlikely to be retrieved on the second and third trial. Similarly, if an item has been retrieved on the first trial, it will receive a boost in strength making it highly likely that it will also be retrieved on the next trials. Suppose 80% of the items are successfully retrieved. These items will then be strengthened as a result of the successful retrieval on the first trial, and again on the second and third trials. The 20% of the items that were not successfully retrieved will not be strengthened on the first trial and most likely also not on the following trials. But this implies that although the mean level of recall stays at about 80%, there may be a huge increase in trace strengths that would not be noticeable when one looks at the observed recall probabilities. As a result, although the mean level of recall stays the same, the effect of the increased strengths on the RP- items may be quite substantial. Figure 8.2b shows the predicted results for the Anderson et al. (2000) experiment that were obtained using the SAM model. Kornell, Bjork, and Garcia (2011) independently developed

a similar explanation to account for the so-called testing effect (for which the comparison between retrieval and study conditions is also crucial).

Interestingly, such a model also provides an explanation for a large number of other results that have been interpreted as evidence for the inhibition account. Several studies have demonstrated that RIF disappears when the retrieval practice is given under stress (Koessler, Engler, Riether, & Kissler, 2009) or when a concurrent secondary task has to be carried out (Román, Soriano, Gómez-Ariza, & Bajo, 2009). These findings fit the inhibition framework under the likely assumption that such manipulations compromise the cognitive control process that is assumed to be responsible for the suppression of the RP- items. However, they also fit the competition model described above if one assumes that these manipulations reduce the amount of additional information that is stored for the RP+ items. Such a reduction will have little effect on RP+ recall (the retrieved items will still be boosted although to a lesser extent) but it will affect RP- recall (the reduced strength will decrease RIF). This explanation also predicts the somewhat surprising result that these manipulations have little effect on RP+ recall even though we know that they affect memory storage.

A number of recent experiments have obtained results that appear to be inconsistent with the retrieval specificity assumption. Raaijmakers and Jakab (2012) showed that RIF may be obtained in tasks using noncompetitive retrieval provided that the noncompetitive task is not too easy (generating the category name with the exemplar as cue is usually too easy) and ensures a focus on the associations to the category cue (see also Jonker & MacLeod, 2012). Similarly, Verde (2013) demonstrated that whether or not RIF is obtained depends on whether the study task focuses on the item features or the category features. In a series of experiments, Verde (2013) showed that a standard study task that focuses mostly on item features will not show RIF (i.e., a replication of the standard finding). However, a study task that emphasizes the category features (e.g., asking participants to judge whether a given category would be the best one to choose for a specific item or whether another category would fit that item better) did lead to a sizable RIF effect.

These findings were replicated by Grundgeiger (2014) and Rupperecht and Bäuml (2016). More importantly, however, these researchers also showed that RIF disappeared when a recognition test was substituted for the standard cued recall test with the category name plus initial letters as cues. These result led them to propose a revised version of the inhibition account in which RIF on a cued recall test is due to a combination of inhibition and competition while RIF on a recognition test is due to inhibition only and is not affected by competitor interference. This seems a sensible hypothesis that solves many of the problems surrounding the inhibition account. I will return to this hypothesis in the final section on the current status on the inhibition account. For the moment, it should just be mentioned that alternative interpretations still seem possible, e.g., the

hypothesis that the observed difference simply reflects a difference in susceptibility to interference effects (i.e., interference effects might need to be stronger to show up on recognition tests).

Strength independence. According to the inhibition account, the extent to which the RP- items show RIF on the final test is assumed to depend only on the extent to which they have been suppressed. It should be independent of the extent to which the RP+ items have strengthened during the retrieval practice phase (Anderson, 2003), provided the final test is not affected by output interference. This assumption is termed strength independence. It contrasts with predictions of competition-based models that assume that the decrease for RP- is caused by the increase in RP+, hence these two should be correlated. Note that this assumption illustrates once more that the inhibition account assumes that competition plays no role (or only a very minor one) at the final recall test.

Tests for this strength independence assumption often involve the same type of experiments as are used to test the retrieval specificity assumption. A comparison is usually made between the standard retrieval practice condition and a condition in which the RP+ items are also strengthened but not in a way that involves competitive retrieval. As before, the inhibition account either predicts no RIF or normal RIF (depending on the nature of the practice task). Inhibition proponents have claimed that a competition-based account should predict RIF in both cases because what matters is that the RP+ items have been strengthened. As we have seen above, the simulation results with the SAM model (Raaijmakers & Jakab, 2012) show that such a conclusion is incorrect since the observed pattern (equal RP+ recall, RIF vs no RIF) can be easily explained using a competition-based model.

More recently, another type of evidence has been proposed for the strength independence assumption. Here, the analysis focuses on the correlation between the strengthening of the RP+ items and the size of the RIF effect. Such analyses have been reported by Hanslmayr, Staudigl, Aslan, and Bäuml (2010), Staudigl, Hanslmayr, and Bäuml (2010), Hulbert, Shivde, and Anderson (2012), and Rupperecht and Bäuml (2016). A meta-analysis carried out across a large number of experiments within the RIF paradigm by Murayama, Miyatsu, Buchli, and Storm (2014) also examined the correlation between these two measures. However, since it is unclear what to make of such correlations across experiments, I will restrict the discussion to experiments in which the increase in mean RP+ recall and the decrease in mean RP- recall were correlated over subjects. The standard outcome of such analyses is that the obtained correlations are small and not significant. The lack of a correlation fits the strength independence principle and appears to be inconsistent with competition based accounts that should predict a positive correlation between the strengthening of RP+ and the magnitude of the RIF effect.

As already mentioned by Raaijmakers and Jakab (2013), however, the lack of such a correlation should be interpreted with care since several factors may mask the correlation. To investigate this issue, Raaijmakers (2016) ran several large-scale simulation studies using a simplified version of the SAM model. The results showed that such a model is unlikely to predict a correlation despite the fact that at the level of the predicted probabilities the correlation is clearly present. The results showed that the lack of a correlation is mostly due to the fact that recall scores are inevitably measured as 1/0 scores (correct/incorrect) and this introduces enough variability to mask the theoretical relation. Hence, such correlations do not provide a good test for the strength independence principle and will not be able to distinguish between inhibitory and non-inhibitory explanations.

Cue independence. As mentioned in the description of the inhibition theory, it is assumed (at least in the version proposed by Anderson, 2003) that it is not the association between the cue and an item that is inhibited but the item representation itself. This implies that a decrease in performance on that item should be observed not just when the item is tested using the category cue it was paired with during study but also when any other cue is used that would normally with some probability activate the item. Anderson and Spellman (1995) devised a rather complicated series of experiments in which they tested this assumption. In these experiments, items were presented as exemplars of one category but could also be considered a member of another category on the list (e.g., strawberry was presented as a member of the category RED with another category FOOD on the same list). They showed that items that were related in this way to one of the RP categories also showed RIF even though they were tested with a different category cue.

Although there are alternative explanations for such a result (see Raaijmakers & Jakab, 2013), these would probably not work for another experiment of Anderson and Spellman (1995, Exp. 2). In this experiment unrelated categories were used but some exemplars from category A and some exemplars from category B were themselves part of a (not presented) category C. What they observed was that if the exemplars from category A were suppressed, there was RIF not only for those exemplars but also for the related exemplars from category B. This is an intriguing although somewhat surprising result. It does not follow from the standard inhibition account (after all, these items would not be activated during the retrieval practice and therefore would not need to be suppressed). To explain these results, Anderson and Spellman (1995) formulated the so-called pattern suppression model. In this model it is assumed that items (exemplars) are represented by features and, when a given feature is suppressed, the suppression spreads to all traces that share that feature. Although the model has not been elaborated in sufficient detail to allow numerical predictions, it is of interest as it generalizes the notion of spreading activation to what might be called spreading inhibition.

It is not clear, however, whether this pattern suppression model would still make the same predictions as the standard inhibition account (see Raaijmakers & Jakab, 2013). Moreover, as shown by Norman, Newman, and Detre (2007) and Raaijmakers and Jakab (2013), it makes a number of predictions that do not seem to be correct. Finally, there is some doubt whether the empirical result obtained by Anderson and Spellman (1995) is in fact reliable. This cross-category RIF effect could not be replicated by Williams and Zacks (2001), Camp, Pecher, and Schmidt (2005) and Rowland, Bates, and DeLosh (2014).

The cue independence prediction does not stand or fall with this specific result. A number of studies have used novel, extralist item-specific cues to test for cue independence. For example, if the target list item was 'guitar', a novel cue item 'musical instrument' might be used. The novel cue item is always chosen in such a way that it is only associated to that specific item and not to any other list item. Both Saunders and MacLeod (2006) and Aslan, Bäuml, and Pastötter (2007) used such a procedure and obtained a sizable cue-independent RIF effect, although the effect was not found by Camp, Pecher, and Schmidt (2007).

There has been an extensive discussion in the literature whether these results might be due to what is called 'covert cuing', the idea that the participants might be covertly retrieving the original category name and using that as a cue, presumably in addition to the extralist cue (see Perfect et al., 2004; Camp, Pecher, & Schmidt, 2005, 2007). However, Hulbert, Shivde, and Anderson (2012), Huddleston and Anderson (2012) and Weller, Anderson, Gómez-Ariza, and Bajo (2013) argued against such an assumption and showed that the use of such covert cues might increase rather than decrease recall. It is not clear why covert cuing would mask RIF unless such a strategy would only be used for the RP- items and not for NRP items. One might argue that it is more difficult to retrieve the RP- items and that that is the reason why the participants resort to such a strategy, but in most experiments the difference between the recall of the RP- and the NRP items is not large, making it unlikely that the strategy would only be used for the RP- items.

Hanczakowski and Mazzoni (2013) carried out a similar experiment in which the additional cues were episodically related to the target items. In their experiments, there were three initial presentations of the list of category-exemplar pairs. On the second and third presentations however an additional unrelated word was presented alongside the category-exemplar pair. They showed that RIF was obtained if the additional cue was presented together with the category cue but RIF was not obtained when the additional cue was presented without the category cue. A similar finding was previously obtained by Perfect et al. (2004). In both cases, there was a failure of cue independence when the target items were tested using item-specific episodically related cues.

Jonker, Seli, and MacLeod (2012) set up an ingenious experiment in which they used the standard retrieval induced forgetting paradigm but the items within a specific category were chosen in such a way that they belonged to one of two subcategories. For example, within the category BIRDS there were two subcategories: 'birds of prey' and 'pet birds'. These subcategories were rather subtle and would not be spontaneously noticed by the participants (as previously shown by Gardiner, Craik, and Birtwistle, 1972). No mention was made of the presence of these subcategories until the final test. The design that they used was such that the practiced items (the RP+ items) were either all from one subcategory (the pure condition) or from both subcategories (the mixed condition). When at the final test only the overall category cue was given (e.g., BIRDS), a standard RIF effect was obtained. However, when the subcategory name was given as a cue, the pattern of results depended on the condition (pure vs mixed). In the mixed condition, standard RIF was obtained but in the pure condition no RIF was found. One interpretation of this result is that in the pure condition there was no interference from RP+ items (that all belonged to a different subcategory). This experiment therefore also failed to support the cue independence assumption.

More positive evidence for cue independence was obtained in an experiment by Veling and Knippenberg (2004). In their experiment, instead of the standard recall test, a lexical decision task was used in which all of the RP+, RP- and NRP items were presented as well as an equal number of novel words and nonwords. Reaction times were higher for the RP- items compared to the NRP items. Butler, Williams, Zacks, and Maki (2001) also examined RIF using a number of other tasks in addition to the standard recall task and did not obtain RIF except in the standard category-cued recall test. However, their results (that seem to conflict with many other experiments that do find RIF with a category-plus-stem cued recall test) might be due to the fact that in their lists there were a large number of item-to-item associations (in addition to the category-item associations). As shown by Goodmon and Anderson (2011), the presence of item-to-item associations between the RP+ and RP- items will lead to a decrease in RIF. Such a result is understandable since a RP- item that is related to a RP+ item may help rather than hinder retrieval of the RP+ item during retrieval practice or may benefit itself from the strengthening of the related RP+ item.

A somewhat controversial line of evidence for cue independence that has been advanced by Anderson and colleagues is the finding that RIF may also be obtained when a recognition test is substituted for the recall test. It is not the finding itself that is controversial but the interpretation that has been given to it. Anderson (2003) and Bäuml (2008) regard this as strong support for the inhibition account but others (Raaijmakers & Jakab, 2013; Verde & Perfect, 2011) have argued that such a result is not necessarily inconsistent with a competition-based account since such accounts may predict interference effects in recognition.

In the next section of this chapter, I will evaluate the relative contributions of the opposing accounts of retrieval induced forgetting and whether these different viewpoints might be reconciled. First, however, I will discuss a number of other experimental paradigms that have been used in the past to investigate inhibition in memory.

The Directed Forgetting paradigm

The oldest experimental paradigm used in the investigation of inhibitory processes is the directed forgetting task (dating back to Brown, 1954). Basically, in this task, the participants are instructed at some point to forget some of the items (suggesting that these will not be tested) and the interest is what happens with these items when they are unexpectedly tested at a final test. The standard finding is that the forget items are indeed recalled less well and the discussion focuses on the question whether this forgetting is due to inhibition or whether there are alternative explanations that might explain the observed decrease in recall.

There are two basic variants of such a task. In the first version, termed the item method, a list of items is presented and each item is followed by either a forget cue or a remember cue (indicating whether the item will be tested later). In the second version, the list method, a series of items is presented and the experimenter announces halfway through the list that the items presented thus far will in fact not be tested and hence might just as well be forgotten. Only the items presented next will be tested. After that, list presentation continues. Hence, in both variants there are to-be-remembered and to-be-forgotten items (R and F items, respectively).

The history of the research on directed forgetting was extensively reviewed by MacLeod (1998). One of the earliest theoretical accounts of directed forgetting was that of Bjork (1972). Bjork emphasized two processes operating in such tasks: (a) the mental separation of R and F items, and (b) the selective rehearsal of the R items. In later years, however, these researchers (see Bjork, 1989) began to emphasize the importance of inhibitory processes for the explanation of the results in this paradigm. A major breakthrough was made when Basden, Basden, and Gargano (1993) pointed out that there was a difference between the item and list methods when memory was tested using recognition. Whereas the item method leads to a decrease in both recall and recognition, the list method leads to a decrease only in recall; no decrease is observed in recognition (although Sahakyan, Waldum, Benjamin, and Bickett, 2009, later showed that under special conditions a decrease may be observed).

To explain this pattern of results, it was proposed that the directed forgetting effects in the item method were due to differential rehearsal (leading to

changes in stored strength that would show up in both recall and recognition) whereas the directed forgetting effects in the list method were due to inhibition. To explain that there are in that case no effects in recognition, it was further assumed that the recognition test (with the inhibited item itself as a cue) would lead to a 'release of inhibition' (Bjork, 1989).

Clearly, the assumption that the inhibition is released indicates that the inhibition that is assumed to be responsible for the decrease in performance in directed forgetting refers to a different process than the inhibition that is assumed to be responsible for RIF, which is longer lasting. This is also indicated by the fact that in list method directed forgetting it is not the forget instruction itself that is responsible for the decrease: it is crucial that there is indeed a second list to be learned (the R items), see Gelfand and Bjork (1985; cited in Bjork, 1989). Such a result is difficult to explain by the inhibition account of Anderson (2003), since studying an item is assumed not to necessitate inhibitory control. Within the inhibition framework, a more likely interpretation is that in the case of directed forgetting it is not the item representation itself but the access to the item that is inhibited. An explanation of this type was proposed by Bäuml (2008).

An alternative explanation for directed forgetting in the list method paradigm has been formulated by Sahakyan and Kelley (2002) within the contextual change framework described in the beginning of this chapter. They proposed that the instruction to forget induces the participants to change the internal context. In the control condition (where no forgetting instructions are given) there will be no such context change. Such a model predicts the standard negative effect on the F-items because the context at test resembles the changed context more than it does the original context (that was stored with the F-items). The model also explains why there will be an enhancement effect for the R-items. These items have been encoded in the changed context and hence the context at test matches the context stored with these items. In effect, for these R-items, the context change leads to a decrease in the effective list length compared to the control condition where all items may be assumed to be encoded in more or less the same context. Thus, the contextual change model can account in principle for both the costs and benefits of directed forgetting. Note that as in the revised inhibition account discussed above, contextual change affects the access to the item representation rather than the item representation itself.

Early support for the view that participants may change the internal context to separate sets of items can be seen in an experiment by Shiffrin (1970) in which participants were presented a series of free recall lists and had to recall the list-before-last. In such a task, having a recall test between the successive lists greatly reduces the interference effect of the number of items on the intervening list (the list presented last) as though it helps to mentally separate the lists. This interpretation was supported by the results of a series of experiments by Jang and Huber (2008).

Within the directed forgetting paradigm, support for the contextual change explanation was first presented by Sahakyan and Kelley (2002). They showed that the costs and benefits of directed forgetting could be mimicked by inducing a context change between the two lists without any forget instructions. Context change was induced by having the participants imagine that they were invisible and then to describe what they would like to do. This created a mental separation between the two lists that was not present in the control condition (where the participants were given standard 'remember' instructions). The results showed that a standard forget instruction and the context change manipulation led to similar results: a decrease in the recall of List 1 items and an increase in the recall of List 2 items. They also showed that mentally reinstating the List 1 context increased recall of List 1 both for the context change condition as well as for the standard forget condition. These results suggest that a standard forget instruction also induces a change in the internal context. Sahakyan and Delaney (2003, 2005) added a second factor, showing that the benefits in directed forgetting (for the R items) are also affected by a change to a more elaborate encoding strategy after the forget cue.

Additional support for this contextual change explanation was obtained by Sahakyan and Goodmon (2010), Mulji and Bodner (2010), Lehman and Malmberg (2011), and Hanczakowski, Pasek, and Zawadzka (2012). Finally, Lehman and Malmberg (2009) implemented the contextual change explanation within the SAM-REM framework and showed that the model accounted well for the existing data.

All in all, then the directed forgetting phenomenon can be explained without recourse to a concept like 'inhibition'. The contextual change account appears to be able to explain most of the results. Interestingly, this account is similar to the 'mental separation' factor proposed by Bjork (1972) in one of the earliest theoretical analyses of the phenomenon.

The Think / No-Think paradigm

Probably the most direct test of the notion that cognitive control may be used to suppress memory traces comes from the Think / No-Think (TNT) paradigm, initially developed by Anderson and Green (2001). As in the retrieval induced forgetting paradigm, there are generally three phases to such an experiment. In the first phase, a list of paired words (usually very weakly associated) is presented. I will refer to the first word as the cue word and the second as the target. Participants are instructed to learn these pairs for a later memory test. After the initial presentation, the cue words are presented and the participants have to generate the corresponding target word. Feedback is given if the response was incorrect (i.e., an anticipation method is used). This continues until some criterion is met (such as 50 or 70% correct). In the second phase, the




Figure 8.3: Final recall for different types of items in the Think / No-Think experiment of Anderson and Green (2001). Panel A shows the results for Experiment 1 in which half of the items had to be suppressed, Panel B shows the results for Experiment 4 in which half of the items were allowed to be recalled but the response should not be said aloud.

think/no-think phase) the cue words are again presented but now for half of the cue words the participants are to give the target word and for the other half they are instructed to try not to think of the corresponding target words, to block these from entering consciousness. To help with this task, often the cue words are presented in either green (think items) or red (no-think items). After the TNT phase, a final test is given in which all cue words are presented and the participants are instructed to generate the corresponding target words, even for those items that they were previously told not to think about.

The Anderson and Green (2001) results are presented in Figure 8.3a. What is of interest here is that the no-think items did less well on the final test compared to the items that were never presented in the TNT phase of the experiment (the baseline condition). A similar result was obtained in a variant of the experiment in which, at the final test, a cue word was presented that was semantically related to the target word but that never been presented before (the independent cue). Anderson and Green (2001) interpreted this finding as evidence for the inhibition of the memory traces for the no-think targets. The fact that the effect also occurred with the independent cues shows that it is not the association between the original cues and the target words that is inhibited but the memory traces of the targets themselves.

The effect has now been replicated a number of times and a meta-analysis reported by Anderson and Huddleston (2012) indicates an average difference between the baseline recall and the recall of the no-think targets of about 8%. Of course, such an analysis may be biased due to non-publication of negative (non-

significant) results but Anderson and Huddleston also report a similar result for all studies (both published and unpublished) carried out in Anderson's lab where such a bias should not be present. Nevertheless there have been frequent failures to obtain a decrease for the supposedly suppressed items, most notably in a series of experiments by Bulevich, Roediger, Balota, and Butler (2006).

Anderson and Huddleston (2012) discuss a number of reasons that might lead to a failure to find the effect. First, not all participants might comply with the instructions so a questionnaire should be used at the end of the experiment to filter out those participants that did not comply (although such a questionnaire was also not used in the original Anderson and Green, 2001, experiment). Second, participants might become less motivated and might not keep up the effort for the whole duration of the experiment. However, these factors do not seem to be able to account for the results of Bulevich et al. (2006) since in these experiments a compliance questionnaire was used and the instructions were copied from Anderson and Green (2001). Not discussed by Anderson and Huddleston (2012) but surely also important is the fact that the effect is small and variable so there should be a substantial probability of non-replication based on statistical considerations, even if the effect is real.

A somewhat peculiar finding that indicates that the paradigm may not be that simple is the result from a control experiment by Anderson and Green (2001) in which the participants were allowed to recall the item but not to say it aloud (the 'withhold' condition). While there was a clear increase in the number of correct responses in the 'respond' condition after 16 trials, performance in the 'withhold' condition did not increase at all (see Figure 8.3b). Normally, one would expect about equal levels of recall in the 'respond' and 'withhold' conditions (e.g., see Putnam & Roediger, 2013).

Part of the problem appears to be that the participants are given very little guidance in how they should perform the task. How does one prevent the target word from coming to mind? Having no control over what strategies participants use might well be another factor that causes variability in the results. Several authors have suggested that participants might use the strategy to think of an alternative word to replace the target word, the so-called substitution strategy. Hertel and Calcaterra (2005) and Lemoult et al. (2010) demonstrated the effectiveness of such a strategy. Del Prete, Hanczakowski, Bajo, and Mazzoni (2015) replicated this result and also observed no forgetting when the conditions were such that it was difficult to generate appropriate substitutes.

Anderson and Huddleston (2012) acknowledge the effectiveness of the substitution strategy but argue that it does not fully explain the effects of thought suppression and that it just increases the size of the effect that would otherwise also be present.

To evaluate this proposal, it is helpful to consider the Hertel and Calcaterra (2005) experiments in more detail. In their experiment, the participants were given a list of adjective-noun pairs in the first phase of the experiment. In the second phase of the experiment (the TNT phase), there were two conditions: the unaided condition (i.e., the standard TNT task) and the aided condition. In the aided condition, the participants were provided with substitute nouns to use during the suppression task. At the final test, they were instructed to recall the original noun but were allowed to give both responses if they recalled both the original noun and the aid. The mean recall at the final test showed no difference between the suppressed and the baseline items for the unaided group but did show a difference for the aided group. In a further analysis, the unaided group was split according to whether the participant reported the spontaneous use of a substitution strategy. In this analysis, it was shown that only those participants that had used a substitution strategy (in this case spontaneously) showed a decrease for the suppressed items. Hertel and Calcaterra also administered a compliance questionnaire and divided the participants into compliant and noncompliant groups (i.e., the noncompliant group occasionally tried to check whether they still knew the original noun). Anderson and Huddleston's critique hinges on the fact that in the unaided condition the compliant participants did show a decrease for the suppressed items and that the overall lack of an effect is due to the mixing of compliant (showing a negative effect) and noncompliant (showing a positive effect) participants. However, that does not imply that the negative effect in the compliant group could not be due to the use of the substitution strategy. Indeed, Hertel and Calcaterra examined the combined effects of compliance and strategy and concluded that "the ability of a substitution strategy to account for success in suppression was not an artefact of whether the participant complied with instructions" (Hertel & Calcaterra, 2005, p. 486).

An alternative explanation for the forgetting effect observed in the TNT task was proposed by Tomlinson, Huber, Rieth, and Davelaar (2009). They showed the same type of forgetting could be obtained if participants were simply asked to quickly press the Enter key instead of trying to suppress the item. Such a result might indicate that the requirement to do something else is responsible for the effect and that suppression may not be crucial.

All in all, it seems fair to conclude that reliable negative effects of suppression can be obtained in the TNT task (although the results tend to be variable) but it cannot be ruled out that these effects are due to other factors such as the use of the substitution strategy. The fact that del Prete et al. (2015) did not obtain a suppression effect when a substitution strategy was unlikely to be used supports such a conclusion.

Other paradigms

There are a few other phenomena where inhibition or suppression has been proposed as a possible mechanism but never gained the popularity that it has as an account for the three paradigms discussed thus far. In this section I will briefly discuss two of these paradigms, output interference and part-list or part-set cuing, without going in too much detail, mainly for historical reasons.

Output interference. In the 1970's, a large number of studies were done to investigate the inhibitory effects of prior retrievals on the probability of recalling a following item. These studies were motivated to a large extent by the landmark experiments of Tulving and Pearlstone (1966) and Slamecka (1968). Tulving and Pearlstone showed that the number of items that can be recalled from a list greatly varies depending on the presence of appropriate retrieval cues. They presented word lists that consisted of exemplars from a number of categories. After the list had been presented, the participants were either given a free recall test or a cued recall test with the category names as cues. Presenting the category cues greatly improved overall recall. An interesting detail of the results was that the advantage of the category cues diminished as the number of items in the category increased. The Slamecka experiment (that I will discuss in more detail below) also appeared to show that presenting additional list items might not always have a positive effect. The explanation of these results focused on the hypothesis that recall might be a self-limiting process (Roediger, 1973, 1974, 1978; Roediger & Neely, 1982).

An instructive experiment is that of Roediger (1978). In this experiment, lists consisting of items from 8 different categories were presented. At recall, some of the participants were given no cues, some all of the 8 category cues and some only 4 of the 8 category cues. It is the latter two conditions that are of interest for the present discussion. In one of these (the inclusion condition), the instructions emphasized to try to recall as many of the items as possible, from all categories but especially from the cued ones. In the other condition (the exclusion condition), the participants were also instructed to recall as many items as possible except from the cued categories. The results showed that the exclusion group recalled many more items from the other, noncued, categories compared to the inclusion group. Hence, recalling items from the cued categories had a negative or inhibitory effect on the recall from the remaining categories. This effect was termed the output interference effect.

Roediger (1974) showed that category cues have a positive effect but that once access to the category has been achieved, presenting additional items from that same category does not help and rather hurts the recall of the remaining items. Roediger (1973) showed that there was a linear decrease in the probability of recall from a category as a function of the number of previously tested categories. Roediger and Schmidt (1980) observed a similar linear decrease for the recall of

paired associates. Both Smith (1971) and Roediger (1973) showed that the output interference effect could be described as a function of the number of items previously recalled.

The standard interpretation of such output interference effects is that recalling an item strengthens its memory representation and this strengthening decreases the probability of sampling subsequent items (due to the fact that sampling is based on the relative strength). Raaijmakers and Shiffrin (1980) showed that all of these results were predicted quite well by the SAM model (that is of course based on exactly these assumptions). Although I have focussed the discussion on recall, the output interference effect also occurs in recognition testing (see Criss, Malmberg, & Shiffrin, 2011; Annis, Malmberg, Criss, & Shiffrin, 2013). For a more extensive review, see Malmberg, Lehman, Annis, Criss, and Shiffrin (2014).

There is some resemblance between this output interference effect and the retrieval induced forgetting effect. In both cases, retrieving a specific item has a negative effect on the later retrieval of other list items. Anderson (2003) and Bäuml (1998) indeed advocated an explanation of output interference in terms of retrieval inhibition. In this view, output interference is not due to an increase in the strength of previously recalled items but to retrieval-induced inhibition of the later items caused by the retrieval of the previously tested items. Bäuml (1998) tested the hypothesis that strong items should suffer more from output interference (since these are more likely to hinder the retrieval of the items tested earlier and hence need not be suppressed). Bäuml (1998), following a suggestion by Anderson et al. (1994), claimed that competition-based accounts of output interference would make exactly the opposite prediction. However, Raaijmakers and Jakab (2013) showed that this prediction does not follow from models based on relative strength for tasks in which mixed lists are used (lists composed of both strong and weak items). They showed that such models will predict a positive linear relation between output interference and the original strength of an item (i.e., strong items should show more output interference). Hence, finding that strong items suffer more from output interference cannot be used as an argument against non-inhibitory models and in favor of the inhibition account.

However, such an account of output interference based on inhibition suffers from a number of problems. First, as mentioned above, output interference appears to depend on the number of previous items successfully recalled rather than just on output position or the number of items tested previously. According to a standard retrieval-induced forgetting explanation, it should not matter whether the previous items are recalled successfully. The only thing that matters is that there is an attempted retrieval (see Storm, Bjork, Bjork & Nestojko, 2006, for a demonstration that retrieval success is not required for RIF to occur). Bäuml (1997), on the other hand, seems to agree that it is the number of successful recalls that matters but such an assumption deviates from the standard inhibition account.

Second, output interference effects occur not only in recall but also in recognition. Although Bäuml (2008) sees this as confirming the cue-independent nature of inhibition, such a conclusion seems inconsistent with the standard assumptions of the inhibition account. In recognition, the item itself is presented and hence there should be no inhibition (there is no interference from other items that has to be resolved by inhibitory control processes).

In conclusion, the inhibitory account of output interference is highly problematic and does not appear to be a viable alternative to the more traditional explanations based on the strengthening of previously recalled items.

Part-list cuing. The part-list cuing paradigm was originally developed by Slamecka (1968; see Raaijmakers & Shiffrin, 1981, for a review). Slamecka's starting point was that most memory theories at that time assumed that in a free recall task in which participants learn a list of unrelated words, an associative network is built up to permit recall of as many of the list items as possible. And indeed, many experiments showed that making it easier to create such interitem associations increased overall recall. Slamecka reasoned that if that is indeed the case, providing the participants with a number of entry points into the network should have a strong effect on the probability of recalling the remaining items in the network. Slamecka therefore provided the participants a number of list items as cues and tested whether this would have the expected positive effect on the recall of the remaining list items (compared to a noncued condition in which no item cues were presented). No such facilitation was observed, not even in lists that consisted of items that should be easily associated (although the mean recall did change quite a bit, but to a similar extent for the cued and the noncued condition). Obviously, the simultaneous presence of both an increase in recall when interitem associations were facilitated and a null effect (and sometimes even a slight negative effect) with part-list cuing presented a problem for theories of memory.

Initially the part-list cuing effect was explained using ideas derived from studies on output interference. For example, Rundus (1973) proposed that the negative effect of part-list cuing could be explained by the strengthening of the cue items when they are read at the beginning of the recall process. Hence, reading or studying the list items that were presented as cues was seen as similar to an implicit recall of these items and the negative effect of part-list cuing became an output interference effect due to the strengthening that followed this implicit recall. Roediger (1973; see also Roediger, 1974, 1978) proposed an hypothesis that could explain both the results from output interference studies and from studies on part-list cuing. Roediger assumed that the cue items would give access to a higher order unit but that, once access to that unit was achieved, additional cue items subsumed under the same higher order unit would decrease the probability of recalling the remaining items from that unit. In these studies, categorized lists were used, and the results indeed showed that presenting one item from a category has a positive effect but any additional item cues from the same

category had a negative effect. In this account, it was assumed that there were no horizontal or inter-item connections, only vertical ones (the connections to the category unit).

During the development of the SAM model, Raaijmakers and Shiffrin (1981) discovered that the model did predict the part-list cuing effect even though the model made extensive use of inter-item associations. Interestingly, simulations showed that within this model the effect was not due to the strengthening after successful recall since the effect was also generated when strengthening was eliminated from the model (see Raaijmakers & Shiffrin, 1981, for an extensive discussion). It turned out that the negative effect was due to a bias in the sampling process when the part-list cues are being used. The SAM explanation is similar to the hypothesis proposed by Roediger (1973) except that the higher order units now refer to subjective, idiosyncratic categories or clusters based on the inter-item associations that are formed during study of the list. The list cues that are presented do give access to these clusters but additional cues belonging to the same cluster will have a negative effect on the probability of recalling additional items from that cluster.

Raaijmakers and Shiffrin (1981) argued that the original expectation of a large positive effect did not take into account that in both the cued and the noncued condition list items would be used as cues, the only difference being that the list cues in the noncued condition were self-generated whereas those in the cued condition were (to a large extent) supplied by the experimenter. When the inter-item associations get stronger, this will therefore increase recall in both conditions. In both cases, the cues will tend to retrieve other list items that are associated to these cues but the subjective clusters that will be accessed by the part-list cues will (necessarily) contain at least one other cue item (namely the list cue that is being used to access the cluster) and therefore will on average contain slightly fewer of the remaining items. However, since it is these items that make up the recall score (both conditions will be scored on the items that are not given for free to the cued group), there will be a slight disadvantage for the cued group, i.e., a slight negative part-list cuing effect.

This explanation for the part-list cuing effect makes a number of predictions that have been tested in experiments (see Raaijmakers & Shiffrin, 1981, and Raaijmakers & Phaf, 1999, for a full discussion). For example, the model predicts that if it becomes difficult to self-generate item cues, the effect will be reversed. Raaijmakers and Phaf (1999) reported an experiment that supports this prediction. In this experiment, the part-list cuing effect was negative with immediate testing but reversed with delayed testing (where recall levels were substantially lower). The effect will also reverse if the cues are given after an initial free recall test (Allen, 1969). Perhaps the most interesting prediction (in the sense that none of the other explanations seem to make that prediction) is that the effect is predicted to be positive when the list consists of a small number of non-

overlapping clusters. Raaijmakers and Phaf (1999) investigated this using a list of items belonging to several categories. It was assumed that with such lists the subjective clusters would correspond to the categorical structure present in the list. They showed that the cues increased the likelihood of recalling items from the cued clusters (as one might have expected). Interestingly, however, they also observed a positive cuing effect for a set of cues that were chosen completely at random (hence no relation to the subjective organization), a result that was also obtained in simulations with the SAM model (see Raaijmakers & Phaf, 1999, Exp 1).

As mentioned before, this part-list cuing effect was initially seen as a manifestation of inhibitory processes in memory, although not in the sense of suppression. Advocates of a suppression-based inhibition account (as in Anderson, 2003) have proposed that the negative effect of part-list cuing is due to factors similar to those operating in the retrieval induced forgetting paradigm. Bäuml and Aslan (2004) assumed that during the initial processing of the cues there is a covert, implicit retrieval of the memory traces of the cue items and this covert retrieval leads to inhibition of associated items.

Such an explanation is unsatisfactory for a number of reasons. First, it is not clear why an implicit retrieval of the cue words would necessitate a suppression of other list items. Presentation of the cues should be similar to a recognition test and should not be affected by interference from other list items. And if there is no interference, there should be no need for suppression. Second, if one assumes that other list items pop up during the processing of the list items, a more natural assumption would be that the participant would recall these items rather than suppress them. After all, those are the items that need to be recalled. Third, this explanation fails to account for a large number of the earlier findings in the part-list cuing literature (Raaijmakers & Jakab, 2013). Fourth, this inhibition explanation predicts that the effect will be mostly on the items associated to the cues since these are the items that would need to be suppressed during the covert retrieval of the list cues. However, Raaijmakers and Shiffrin (1981) already mentioned that such a prediction is incorrect. In addition, Raaijmakers and Phaf (1999) showed that the cuing leads instead to an increase in the recall of the items associated to the cues but to a decrease in the recall of the other items, those that are not associated to the cues. Such a result is exactly the opposite of what is predicted by the inhibition account.

In sum, the inhibition account does not appear to be able to explain the part-list cuing effect. As shown by the simulations of the SAM model, this effect that was originally seen as quite paradoxical turns out to be completely consistent with a model that assumes that part-list cues do what cues always do: increase the likelihood of retrieving items associated to the cues.

Resolving the problems

The analysis in this chapter of the inhibition account of forgetting portrays a rather mixed picture of the current status of the proposal that inhibition underlies much of the forgetting observed in human memory. While this theory provides a viable account of the forgetting observed in the retrieval induced forgetting paradigm and the Think / No-Think paradigm, it does not provide a satisfactory account for the results in several other paradigms that at one time or another were thought to be manifestations of similar inhibitory processes. Moreover, even where the theory remains the dominant explanation, alternative accounts have not been ruled out. In this section, I will discuss some of the persisting problems with the inhibition account and will propose a direction for future theoretical development.

Consistency of assumptions

One of the most conspicuous problems that have not yet been resolved in a satisfactory way concerns the assumption that a recall test with item-specific cues will not be affected by competition from other list items. The only negative effect that is allowed in the account given by Anderson is that competition might affect the time to retrieve the target item, not the probability of retrieving it. According to Anderson (2003, p. 416):

"By this view, memory retrieval presents a special case of a broad class of situations that recruit executive control processes; it is the executive control mechanism that overcomes interference—inhibition—that causes us to forget, not the competition itself. This view departs from the common assumption that forgetting is a passive side effect of the ever-changing structure of memory. The mere storage of interfering traces is not what causes memories to grow less accessible with time. Rather, forgetting, whether incidental or intentional, is produced as a response to interference caused by activated competitors in memory."

While this assumption is rarely explicitly stated, it should be evident that it underlies several of the properties ascribed to the inhibition account. For example, the property of strength independence (Anderson, 2003) would not hold if competition was allowed to play a role because the strength of the RP+ items would matter, contrary to the strength independence principle.

There are two main problems with such an assumption. The first is that the assumption that competition affects only the response time and not the likelihood of a correct response does not appear very likely and would be hard to implement

in any model of memory retrieval. Normally, one would expect that if a factor delays the retrieval of a target, it would also affect the likelihood of ultimate success. Second, the theory is based on the idea that during retrieval practice a cognitive control process is needed to resolve the interference from related items. In a standard retrieval induced forgetting paradigm, during the retrieval practice phase, the category name plus the two initial letters of the target item are presented. At the final test, the category name plus usually the first letter of the target are presented. If the latter test is free from the effects of competition, it should be the case that the retrieval practice trials should also be free of competition. This would lead to the paradoxical result that there would be no need for inhibition since there is no interference to be resolved.

Advocates of the inhibition account are of course aware of this problem, yet there does not appear to be a clear resolution. For example, Schilling, Storm, and Anderson (2014, p. 360) write: “Adding item-specific stem cues, therefore, should reduce (though not eliminate) blocking from Rp+ items during the retrieval of Rp- items at final test.” Similarly, Storm and Levy (2012, p. 829) write: “Importantly, proponents of the inhibitory account do not claim that noninhibitory mechanisms cannot cause retrieval induced forgetting. In fact, response competition from strong items and the resultant blocking of weaker items is precisely the situation that purportedly triggers the need for inhibition. And it seems possible that both inhibition and interference contribute to some degree to all demonstrations of retrieval-induced forgetting.”

This problem becomes particularly clear in the treatment of the so-called ‘correlated costs and benefits problem’ (Anderson & Levy, 2007; Schilling et al., 2014). The issue here is that it is often found that special populations that are known to have problems with cognitive control (e.g., ADHD patients) do not show less retrieval induced forgetting than normal groups. Anderson and Levy correctly pointed out that this might be due to the fact that these patient groups indeed lack the ability to inhibit the related items during the retrieval practice phase but that this same inability will make it difficult for them to retrieve the RP- items on the final test due to the interference from the (now very strong) RP+ items. According to Schilling et al. (2014), this should occur especially in tests that do not use item-specific cues, such as a standard category-cued recall test. Evidence consistent with such an analysis was obtained by Storm and White (2010) in which normal RIF was obtained with ADHD patients on a category-cued recall test but there was no RIF on a test using item-specific cues.

However, since the recall during the retrieval practice is also tested using item-specific cues, there should also be no problem for the ADHD group during the retrieval practice phase, i.e., there should be no interference to overcome. Storm and White (2010) indeed observed no difference between the ADHD and non-ADHD groups in their performance during the retrieval practice phase. Clearly, if the lack of inhibitory control ability has no effect on performance in the

retrieval practice phase, it becomes unclear what purpose inhibition is supposed to serve.

In the original discussion of the correlated costs and benefits problem (Anderson & Levy, 2007), this problem does not arise and the retrieval practice trials and the final test trial are treated similarly. Anderson and Levy assume that inhibitory control serves the purpose to resolve the interference from related items. Participants with intact inhibitory capacity will inhibit the RP- items during the retrieval practice. When these items are later tested, their strength will be reduced and, because the inhibitory control is effective, they will not show much additional decrease due to the strengthened RP+ items. In contrast, participants whose inhibitory control is not very effective will not inhibit the RP- items much during the retrieval practice. For the same reason, these participants will have problems resolving the interference by the strengthened RP+ items on the final test and hence should show a decrease in RP- recall, despite the fact that these items have not been inhibited. Although this analysis would seem to be more consistent with the general inhibitory framework, it does predict a decrease for the RP- items, even on a final test that makes use of item-specific cues, a result that according to Schilling et al. (2014) is not observed.

In recent years, a number of researchers (Grundgeiger, 2014; Rupperecht & Bäuml, 2016) have proposed to relax the assumption that item-specific cues eliminate competition (at least in recall). However, as yet there has been no coherent account based on inhibition that reconciles this assumption with other assumptions of the inhibition account (e.g., strength independence). In the next section, I will outline a reformulation of the inhibition account that might solve these problems.

Steps toward a reformulation

A possible reformulation of the inhibition account should take as its starting point the assumption that an inhibitory control process is invoked whenever there are alternative memory traces that make it difficult to retrieve the target trace. This will be the case whenever the combined strength of the competitor traces is high *relative to that of the target trace*. Thus, low-strength competitors may still be interfering if the target strength is low. In the standard version of the inhibition account as proposed by Anderson (2003), the effect of the original strength of the RP+ items is largely ignored but it would seem to make sense even in an inhibitory framework that there should be less need for inhibitory control to overcome interference if there is little interference to begin with.

The extent to which a competitor item interferes with the retrieval of the target item should also depend on the extent to which the available retrieval cues focus on the target item. That is, there should be less competition with item-specific cues and perhaps no competition with recognition testing. However, even with item-specific cues such as the initial one or two letters, there will still be competition from other memory traces (see Grundgeiger, 2014; Rupprecht & Bäuml, 2016), both during the retrieval practice trials and on the final test.

Finally, some decisions would have to be made about the inhibitory process itself. I see two (not necessarily exclusive) possibilities. First, the strength of all competing traces might be reduced, perhaps in proportion to their current strength levels. Note that this inhibition would take place regardless of whether a specific competitor trace has been activated (sampled) or not. Alternatively, only the strength of the specific competitor item that was activated will be reduced; all other competitors stay at their current value.

The analysis of such a hybrid model that combines features of inhibitory and non-inhibitory models will be pretty complicated, not because the model itself is complicated but because of the intrinsic dependencies during the retrieval process. In this respect, the model is similar to models such as SAM and therefore the analysis will probably require Monte Carlo simulations.

To give an example, suppose that on the initial retrieval practice trial, the target RP+ item is successfully retrieved and recalled. The strength of that item will then receive a considerable increase in strength (the practice effect). As a result, this strengthened RP+ item will then increase the interference for the next item that is tested. Hence, RP+ items tested later will suffer more from interference by other items, will have a lower recall and, as a corollary, the inhibitory effects will become larger as testing continues and as more and more RP+ items are strengthened. Usually each RP+ item is presented three times during the retrieval practice phase. Whether there will be a need for the inhibitory control process on the later trials will depend on whether the increase in strength after the first successful recall is large enough to offset the increased interference by the other items (including the other RP+ items). If the increase is large enough, there will be less need for inhibition on the later practice trials for the same target item so most of the inhibition should occur during the initial retrieval practice trials.

At the final test, the RP- items will show reduced recall both because they have been repeatedly inhibited but also because of the interference by the strengthened RP+ items. Since the inhibitory control process will still be operating, there will also be substantial inhibition, especially for the strong RP+ items (assuming that the inhibition affects strong competitors more than weak ones). If the RP- item is nevertheless retrieved it will receive a boost in strength (just as the RP+ items did during the retrieval practice phase). Hence, if all of the

RP- items are tested before all of the RP+ items (not unusual in these experiments) performance on the RP+ items will be severely affected both because of the additional strengthening of the RP- items and because of the inhibition that took place during the retrieval of these RP- items.

Clearly, the resulting model will be quite complex and it will often be difficult to anticipate what the model will predict in a given situation. However, in this respect the present model would not be different from other contemporary memory models.

Extending the inhibition account

The classical interference theory of forgetting (Postman, 1961; Postman & Underwood, 1973) distinguished between two types of interference: proactive and retroactive interference. Of these, retroactive interference has been investigated the most and also appears to be a more reliable phenomenon. Proactive interference is demonstrated when the recall of material learned later (say List 2) is affected by the prior learning of related information (List 1). Proactive interference is most easily obtained when the two lists are learned in close proximity and the interval between the learning of the second list and the final test is relatively long (say 24 hours).

The occurrence of proactive interference on tests that supposedly eliminate the effects of competition was a problem for the traditional interference theory since unlearning could not play a role, and if both unlearning and competition are eliminated there were no more factors that could account for the proactive interference that was observed. The inhibition account also appears to have problems explaining proactive interference effects. During the presentation of List 2, inhibition could occur for List 1, but there should be no need to inhibit the List 2 items themselves. Hence, proactive interference cannot be due to inhibition. The only factor that remains is competition and there is no reason to expect competition to have an effect if the test allows the recall of both the List 1 and the List 2 items. Hence these results also point to the need to incorporate competition in a revised inhibition model as described above.

Another finding that might require a revision of the standard inhibition account is the phenomenon of retroactive facilitation. Many experiments have shown that the administration of certain pharmacological agents such as benzodiazepines affects the storage or consolidation of new information in memory. If such drugs are given in a two-list interference design before the learning of the second list, the results on a final test given later when the drug is no longer active show that recall of the second list is impaired (as it should be since the drug interferes with memory storage). What is of interest here is that the decrease in recall of L2 is accompanied by an increase in the recall of L1

(compared to a placebo control group). Thus, the second list leads to facilitation in the recall of the first list, hence the term *retroactive facilitation*. This finding is typically explained by the assumption that the decreased strength of L2 makes it less interfering and hence promotes increased recall of L1 (see Hinrichs, Ghoneim, & Mewaldt, 1984).

What the inhibition hypothesis would predict for such experiments will depend on what is assumed with respect to the effect that these drugs have on the ability for inhibitory control. One possibility would be that these drugs also affect the ability for inhibitory control. This does not seem likely since I have found no indication in the literature that benzodiazepines affect cognitive control or the brain areas assumed to be involved in cognitive control. Nevertheless such a hypothesis cannot be completely ruled out at the moment. In this case, one would predict that the drugs will not affect the strength of the L1 items (that normally would be suppressed during the L2 learning). Hence, there should not be a difference between the interference and the control conditions, that is, no retroactive facilitation would be predicted. A second possibility is that these drugs do not affect the ability for cognitive control. In that case, there should be an especially strong inhibition of the L1 items since even after a number of study trials on L2, the L2 items will not be learned and hence the L1 items will continue to compete during the study trials of L2. Hence, it seems difficult for current versions of the inhibition account to predict this phenomenon of retroactive facilitation, unless (again) some effect of competition is allowed.

Conclusion

The renewed attention to the idea that forgetting might be at least partly caused by a process of inhibition has had (and still has) a great impact on research into the causes of human forgetting. It has led to several new experimental paradigms (retrieval induced forgetting and the think/no-think paradigm) and the research in these new paradigms has led to results that are a challenge to most traditional theories of memory. Most of the research has tried to establish properties that supposedly “uniquely support the inhibitory hypothesis” (Anderson, 2003). However, some of these results can in fact be explained quite well by alternative hypotheses that do not rely on inhibition. In addition, counter examples have been given of results that do not seem to be consistent with the properties that supposedly uniquely characterize the inhibition hypothesis. What these results show is that there is a need for a more formalized, computational model of the inhibition account that would make it possible to determine unambiguously what the model predicts and that would be applied in a consistent

manner from one experiment to the next. Such a model might then be compared to the models based on competition/interference that have already been shown to give a successful account for many of the phenomena discussed in this chapter.

Acknowledgements

I would like to thank Ken Malmberg, David Huber, John Wixted, and especially Colin MacLeod for their helpful comments on the initial version of this chapter.

References

- Allen, M. M. (1969). Cueing and retrieval in free recall. *Journal of Experimental Psychology*, *81*, 29-35.
- Anderson, J.R. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior*, *22*, 261-295.
- Anderson, J. R., Bothell, D., Lebiere, C. and Matessa, M. 1998. An integrated theory of list memory. *Journal of Memory and Language*, *38*, 341-380.
- Anderson, M. C. (2001). Active forgetting: evidence for functional inhibition as a source of memory failure. *Journal of Aggression, Maltreatment, & Trauma*, *4*, 185–210.
- Anderson, M. C. (2003). Rethinking interference theory: Executive control and the mechanisms of forgetting. *Journal of Memory and Language*, *49*, 415-445.
- Anderson, M. C. (2005). The role of inhibitory control in forgetting unwanted memories: A consideration of three methods. In N. Ohta, C. M. MacLeod, & B. Utzl (Eds.), *Dynamic cognitive processes* (pp. 159-189). Tokyo: Springer-Verlag
- Anderson, M. C. (2006). Repression: A cognitive neuroscience approach. In M. Mancina (ed.), *Psychoanalysis and neuroscience* (pp. 327-349). Milan: Springer-Verlag.
- Anderson, M. C., & Bell, T. (2001). Forgetting our facts: The role of inhibitory processes in the loss of propositional knowledge. *Journal of Experimental Psychology: General*, *130*, 544-570.
- Anderson, M. C., Bjork, E. L., & Bjork, R. A. (2000). Retrieval-induced forgetting: Evidence for a recall-specific mechanism. *Psychonomic Bulletin & Review*, *7*, 522-530.
- 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 & Cognition*, *20*, 1063-1087.
- Anderson, M.C., & Green, C. (2001). Suppressing unwanted memories by executive control. *Nature*, *410*, 366-369.
- Anderson, M. C., & Huddleston, E. (2012). Towards a cognitive and neurobiological model of motivated forgetting. In R. F. Belli (Ed.), *True and false recovered memories: Toward a reconciliation of the debate* (Vol. 58, pp. 53–120). New York: Springer. Nebraska Symposium on Motivation.
- Anderson, M. C., & Levy, B. J. (2007). Theoretical issues in inhibition: Insights from research on human memory. In D. Gorfein & C. MacLeod (Eds.), *Inhibition in cognition* (pp. 81_102). Washington, DC: American Psychological Association.
- 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.

- Annis, J., Lenes, J. G., Westfall, H. A., Criss, A. H., & Malmberg, K. J. (2015). The list-length effect does not discriminate between models of recognition memory. *Journal of Memory and Language*, *85*, 27-41.
- Annis, J., Malmberg, K. J., Criss, A. H., & Shiffrin, R. M. (2013). Sources of interference in recognition testing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*, 1365.
- Aslan, A., Bäuml, K.-H., & Pastötter, B. (2007). No inhibitory deficit in older adults' episodic memory. *Psychological Science*, *18*, 72-78.
- Basden, B. H., Basden, D. R., & Gargano, G. J. (1993). Directed forgetting in implicit and explicit memory tests: A comparison of methods. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 603-616.
- Bäuml, K.-H. (1997). The list-strength effect: Strength-dependent competition or suppression. *Psychonomic Bulletin & Review*, *4*, 260-264.
- Bäuml, K.-H. (1998). Strong items get suppressed, weak items do not: The role of item strength in output interference. *Psychonomic Bulletin & Review*, *5*, 459-463.
- Bäuml, K.-H. (2002). Semantic recall can cause episodic forgetting. *Psychological Science*, *13*, 356-360.
- Bäuml, K.-H. (2008). Inhibitory processes. In H. L. Roediger, III (Ed.), *Cognitive psychology of memory. Vol. 2 of Learning and memory: A comprehensive reference*, 4 vols. (J. Byrne Editor), pp. 195-220. Oxford: Elsevier.
- Bäuml, K.-H., & Aslan, A. (2004). Part-list cuing as instructed retrieval inhibition. *Memory & Cognition*, *32*, 610-617.
- Bjork, R. A. (1972). Theoretical implications of directed forgetting. In A.W. Melton & E. Martin (Eds.), *Coding processes in human memory*. (Pp. 217-235.). Washington, DC: Winston.
- 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.
- Brown, J. (1954). The nature of set-to-learn and of intra-material interference in immediate memory. *Quarterly Journal of Experimental Psychology*, *6*, 141-148.
- Bulevich, J. B., Roediger, H. L., III, Balota, D. A., & Butler, A. C. (2006). Failures to find suppression of episodic memories in the think/no-think paradigm. *Memory & Cognition*, *34*, 1569-1577.
- Butler, K. M., Williams, C. C., Zacks, R. T., & Maki, R. H. (2001). A limit on retrieval-induced forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 1314-1319.
- Camp, G., Pecher, D., & Schmidt, H. G. (2005). Retrieval-induced forgetting in implicit memory tests: The role of test awareness. *Psychonomic Bulletin & Review*, *12*, 490-494.
- Camp, G., Pecher, D., & Schmidt, H.G. (2007). No retrieval-induced forgetting using item-specific independent cues: Evidence against a general inhibitory account. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 950-958.

- Camp, G., Pecher, D., Schmidt, H.G., & Zeelenberg, R. (2009). Are independent probes truly independent? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*, 934-942.
- Chan, J. C., Erdman, M. R., & Davis, S. D. (2015). Retrieval induces forgetting, but only when nontested items compete for retrieval: Implication for interference, inhibition, and context reinstatement. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*, 1298-1315.
- Ciranni, M. A., & Shimamura, A. P. (1999). Retrieval-induced forgetting in episodic memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 1403-1414.
- Criss, A. H., Malmberg, K. J., & Shiffrin, R. M. (2011). Output interference in recognition memory. *Journal of Memory and Language*, *64*, 316-326.
- Crowder, R.G. (1976). *Principles of learning and memory*. Hillsdale, N.J.: Erlbaum.
- del Prete, F., Hanczakowski, M., Bajo, M. T., & Mazzoni, G. (2015). Inhibitory effects of thought substitution in the think/no-think task: Evidence from independent cues. *Memory*, *23*, 507-517.
- Gardiner, J.M., Craik, F.I.M. & Birtwistle, J. (1972). Retrieval cues and release from proactive inhibition. *Journal of Verbal Learning and Verbal Behavior*, *11*, 778-783.
- Gelfand, H., & Bjork, R. A. (1985). On the locus of retrieval inhibition in directed forgetting. Paper presented at the 26th annual meeting of the Psychonomic Society, Boston MA.
- Gómez-Ariza, C. J., Fernandez, A., & Bajo, M. T. (2012). Incidental retrieval-induced forgetting of location information. *Psychonomic Bulletin & Review*, *19*, 483-489.
- Goodmon, L. B., & Anderson, M. C. (2011). Semantic integration as a boundary condition on inhibitory processes in episodic retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*, 416
- Grundgeiger, T. (2014). Noncompetitive retrieval practice causes retrieval-induced forgetting in cued recall but not in recognition. *Memory & Cognition*, *42*, 400-408.
- Hanczakowski, M., & Mazzoni, G. (2013). Contextual match and cue-independence of retrieval-induced forgetting: Testing the prediction of the model by Norman, Newman, and Detre (2007). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*, 953-958.
- Hanczakowski, M., Pasek, T., & Zawadzka, K. (2012). Context-dependent impairment of recollection in list-method directed forgetting. *Memory*, *20*, 758-770.
- Hanslmayr, S., Staudigl, T., Aslan, A., & Bäuml, K. H. (2010). Theta oscillations predict the detrimental effects of memory retrieval. *Cognitive, Affective, & Behavioral Neuroscience*, *10*, 329-338.
- Hertel, P. T., & Calcaterra, G. (2005). Intentional forgetting benefits from thought substitution. *Psychonomic Bulletin & Review*, *12*, 484-489.
- Hinrichs, J. V., Ghoneim, M. M., & Mewaldt, S. P. (1984). Diazepam and memory: retrograde facilitation produced by interference reduction. *Psychopharmacology*, *84*, 158-62.

- Huddleston, E., & Anderson, M. C. (2012). Reassessing critiques of the independent probe method for studying inhibition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*, 1408-1418.
- Hulbert, J. C., Shivde, G., & Anderson, M. C. (2012). Evidence against associative blocking as a cause of cue-independent retrieval-induced forgetting. *Experimental Psychology*, *59*, 11-21.
- Jakab, E. & Raaijmakers, J. G. W. (2009). The role of item strength in retrieval induced forgetting. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *35*, 607-617.
- Jang, Y., & Huber, D. E. (2008). Context retrieval and context change in free recall: Recalling from long-term memory drives list isolation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*, 112–127.
- Johansson, M., Aslan, A., Bäuml, K.H., Gäbel, A. & Mecklinger, A. (2007). When remembering causes forgetting: Electrophysiological correlates of retrieval-induced forgetting. *Cerebral Cortex*, *17*, 1335-1341.
- Jonker, T. R., & MacLeod, C. M. (2012). Retrieval-induced forgetting: Testing the competition assumption of inhibition theory. *Canadian Journal of Experimental Psychology*, *66*, 204-211.
- Jonker, T. R., Seli, P., & MacLeod, C. M. (2012). Less we forget: Retrieval cues and release from retrieval-induced forgetting. *Memory & Cognition*, *40*, 1236–1245.
- Jonker, T. R., Seli, P., & MacLeod, C. M. (2013). Putting retrieval-induced forgetting in context: an inhibition-free, context-based account. *Psychological Review*, *120*, 852-872.
- Koessler, S., Engler, H., Riether, C., & Kissler, J. (2009). No retrieval-induced forgetting under stress. *Psychological Science*, *20*, 1356-1363.
- Kornell, N., Bjork, R. A., Garcia, M. A. (2011). Why tests appear to prevent forgetting: A distribution-based bifurcation model. *Journal of Memory and Language*, *65*, 85-97.
- Lehman, M., & Malmberg, K. J. (2009). A global theory of remembering and forgetting from multiple lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*, 970–988.
- Lemoult, J., Hertel, P.T., & Joormann, J. (2010), Training the forgetting of negative words: The role of direct suppression and the relation to stress reactivity. *Applied Cognitive Psychology*, *24*, 365–375.
- Loftus, E.F. & Loftus, G.R. (1980). On the permanence of stored information in the human brain. *American Psychologist*, *35*, 409-420.
- MacLeod, C. M. (1998). Directed forgetting. In J. M. Golding & C. M. MacLeod (Eds.), *Intentional forgetting: Interdisciplinary approaches* (pp. 1–57). Mahwah, NJ: Erlbaum.
- Malmberg, K. J., Lehman, M., Annis, J., Criss, A. H., & Shiffrin, R. M. (2014). Consequences of testing memory. In B. H. Ross (Ed.), *Psychology of learning and motivation*, Vol. 61. (Pp. 285-313). USA: Academic Press.

- Malmberg, K. J., Lehman, M., & Sahakyan, L. (2006). On the cost and benefit of taking it out of context: Modeling the inhibition associated with directed forgetting. In Proceedings of the 28th Meeting of the Cognitive Science Society (pp. 549-554).
- Melton, A.W., & Irwin, J. McQ. (1940). The influence of degree of interpolated learning on retroactive inhibition and the overt transfer of specific responses. *The American Journal of Psychology*, *53*, 173-203.
- Mensink, G. J. M., & Raaijmakers, J.G.W. (1988). A model for interference and forgetting. *Psychological Review*, *95*, 434-455.
- Mensink, G.J.M. & Raaijmakers, J.G.W. (1989). A model of contextual fluctuation. *Journal of Mathematical Psychology*, *33*, 172-186.
- Mulji, R., & Bodner, G.E. (2010). Wiping out memories: New support for a mental context change account of directed forgetting. *Memory*, *18*, 763-773.
- Murayama, K., Miyatsu, T., Buchli, D., & Storm, B. C. (2014). Forgetting as a consequence of retrieval: A meta-analytic review of retrieval-induced forgetting. *Psychological Bulletin*, *140*, 1383-1409.
- Norman, K. A., Newman, E. L., & Detre, G. (2007). A neural network model of retrieval-induced forgetting. *Psychological Review*, *114*, 887-953.
- Parker, E. S., Cahill, L., & McGaugh, J. L. (2006). A case of unusual autobiographical remembering. *Neurocase*, *12*, 35-49.
- Perfect, T. J., Stark, L., Tree, J. J., Moulin, C. J. A., Ahmed, L., & Hutter, R. (2004). Transfer appropriate forgetting: The cue-dependent nature of retrieval-induced forgetting. *Journal of Memory and Language*, *51*, 399-417.
- Postman, L. (1961). The present status of interference theory. In C.N. Cofer (Ed.), *Verbal learning and verbal behavior*. Pp. 152-179. New York: McGraw-Hill.
- Postman, L. & Underwood, B.J. (1973). Critical issues in interference theory. *Memory and Cognition*, *1*, 19-40.
- Putnam, A. L., & Roediger III, H. L. (2013). Does response mode affect amount recalled or the magnitude of the testing effect? *Memory & Cognition*, *41*, 36-48.
- Raaijmakers, J. G. W. (2016). On testing the strength independence assumption in retrieval-induced forgetting. *Psychonomic Bulletin & Review*, *23*, 1374-1381.
- Raaijmakers, J. G. W., & Jakab, E. (2012). Retrieval-induced forgetting without competition: Testing the retrieval specificity assumption of the inhibition theory. *Memory & Cognition*, *40*, 19-27.
- Raaijmakers, J. G. W., & Jakab, E. (2013). Rethinking inhibition: On the problematic status of the inhibition theory for forgetting. *Journal of Memory and Language*, *68*, 98-122.
- Raaijmakers, J. G. W & Phaf, R. H. (1999). Part-list cuing revisited: A test of the SAM explanation. In C. Izawa (Ed.), *On memory: Evolution, progress and reflection on the 30th anniversary of the Atkinson-Shiffrin model*. (Pp 87-104). Mahwah, N.J.: Lawrence Erlbaum Associates.

- Raaijmakers, J. G. W. & Shiffrin, R. M. (1980). SAM: A theory of probabilistic search of associative memory. In G.H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory*. (Vol. 14). (pp. 207-262). New York: Academic Press.
- Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, 88, 93-134.
- Roediger, H. L. (1973). Inhibition in recall from cueing with recall targets. *Journal of Verbal Learning and Verbal Behavior*, 12, 644-657.
- Roediger, H. L. (1974). Inhibiting effects of recall. *Memory & Cognition*, 2, 261-296.
- Roediger, H. L. (1978). Recall as a self-limiting process. *Memory & Cognition*, 6, 54-63.
- Roediger, H.L. & Neely, J.H. (1982). Retrieval blocks in episodic and semantic memory. *Canadian Journal of Psychology*, 36, 231-242.
- Roediger, H. L. & Schmidt, S. R. (1980). Output interference in the recall of categorized and paired-associate lists. *Journal of Experimental Psychology: Human Learning and Memory*, 6, 91-105.
- Román, P., Soriano, M. F., Gómez-Ariza, C. J., & Bajo, M. T. (2009). Retrieval-induced forgetting and executive control. *Psychological Science*, 20, 1053-1058.
- Rowland, C. A., Bates, L. E., & DeLosh, E. L. (2014). On the reliability of retrieval-induced forgetting. *Frontiers in psychology*, 5, Article 1343.
- Rundus, D. (1973). Negative effects of using list items as recall cues. *Journal of Verbal Learning and Verbal Behavior*, 12, 43-50.
- Rupprecht, J., & Bäuml, K. H. T. (2016). Retrieval-induced forgetting in item recognition: Retrieval specificity revisited. *Journal of Memory and Language*, 86, 97-118.
- Sahakyan, L., & Delaney, P. F. (2003). Can encoding differences explain the benefits of directed forgetting in the list-method paradigm? *Journal of Memory and Language*, 48, 195–201.
- Sahakyan, L., & Delaney, P. F. (2005). Directed forgetting in incidental learning and recognition testing: Support for a two-factor account. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 789–801.
- Sahakyan, L., & Goodmon, L. B. (2010). Theoretical implications of extralist probes for directed forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36, 920–937.
- Sahakyan, L., & Kelley, C. (2002). A contextual change account of the directed forgetting effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 1064–1072.
- Sahakyan, L., Waldum, E. R., Benjamin, A. S., & Bickett, S. P. (2009). Where is the forgetting with list-method directed forgetting in recognition? *Memory & Cognition*, 37, 464–476.
- Saunders, J., & MacLeod, M. D. (2006). Can inhibition resolve retrieval competition through the control of spreading activation? *Memory & Cognition*, 34, 307–322.
- Schilling, C. J., Storm, B. C., & Anderson, M. C. (2014). Examining the costs and benefits of inhibition in memory retrieval. *Cognition*, 133, 358-370.
- Shiffrin, R.M. (1970). Forgetting: Trace erosion or retrieval failure? *Science*, 168, 1601-1603.

- Shiffrin, R. M., Ratcliff, R., & Clark, S. (1990). The list-strength effect: II. Theoretical mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 179–195.
- Shiffrin, R. M., and Steyvers, M. (1997). A model for recognition memory: REM: Retrieving effectively from memory. *Psychonomic Bulletin & Review*, *4*, 145-166.
- Shivde, G., & Anderson, M. C. (2001). The role of inhibition in meaning selection: Insights from retrieval-induced forgetting. In D. S. Gorfein (Ed.), *On the consequences of meaning selection: Perspectives on resolving lexical ambiguity* (pp. 175–190). Washington, DC: American Psychological Association.
- Slamecka, N. J. (1968). An examination of trace storage in free recall. *Journal of Experimental Psychology*, *76*, 504-513.
- Smith, A.D. (1971). Output interference and organized recall from long-term memory. *Journal of Verbal Learning and Verbal Behavior*, *10*, 400-408.
- Staudigl, T., Hanslmayr, S., & Bäuml, K.-H. T. (2010). Theta oscillations reflect the dynamics of interference in episodic memory retrieval. *The Journal of Neuroscience*, *30*, 11356-11362.
- Storm, B. C. (2010). Retrieval-induced forgetting and the resolution of competition. In A. Benjamin (Ed.), *Successful remembering and successful forgetting: A Festschrift in honor of Robert A. Bjork*. (pp. 89-105). London: Psychology Press.
- Storm, B. C., Angello, G., Buchli, D. R., Koppel, R. H., Little, J. L., & Nestojko, J. F. (2015). A review of retrieval-induced forgetting in the contexts of learning, eyewitness memory, social cognition, autobiographical memory, and creative cognition. In B. H. Ross (Ed.), *The psychology of learning and motivation*. (Vol 62, pp. 141-194). New York: Academic Press.
- 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. (2012). On the durability of retrieval-induced forgetting. *Journal of Cognitive Psychology*, *24*, 617-629.
- 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 & Review*, *13*, 1023-1027.
- Storm, B. C., & Levy, (2012). A progress report on the inhibitory account of retrieval-induced forgetting. *Memory & Cognition*, *40*, 827-843.
- Storm, B. C., & White, H. A. (2010). ADHD and retrieval-induced forgetting: Evidence for a deficit in the inhibitory control of memory. *Memory*, *18*, 265–271.
- Tempel, T., & Frings, C. (2015). Interference in episodic memory: Retrieval-induced forgetting of unknown words. *Psychological research*, *79*, 795-800.
- Tomlinson, T. D., Huber, D. E., Rieth, C. A., & Davelaar, E. J. (2009). An interference account of cue-independent forgetting in the no-think paradigm. *Proceedings of the National Academy of Sciences*, *106*, 15588-15593.

- Tulving, E. & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, 5, 381-391.
- Veling, H., & Van Knippenberg, A. (2004). Remembering can cause inhibition: Retrieval-induced inhibition as cue independent process. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 315–318.
- Verbruggen, F., & Logan, G. D. (2008). Response inhibition in the stop-signal paradigm. *Trends in cognitive sciences*, 12, 418-424.
- Verde, M. F. (2012). Retrieval-induced forgetting and inhibition: A critical review. In B. H. Ross (Ed.), *Psychology of learning and motivation*, Vol. 56. (Pp.47-80). USA: Academic Press.
- Verde, M. F. (2013). Retrieval-induced forgetting in recall: Competitor interference revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39, 1433.
- Verde, M. F., & Perfect, T. J. (2011). Retrieval-induced forgetting in recognition is absent under time pressure. *Psychonomic Bulletin & Review*, 18, 1166–1171.
- Weller, P. D., Anderson, M. C., Gómez-Ariza, C. J., & Bajo, M. T. (2013). On the status of cue independence as a criterion for memory inhibition: Evidence against the covert blocking hypothesis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39, 1232-1245.
- Williams, C. C., & Zacks, R. T. (2001). Is retrieval-induced forgetting an inhibitory process? *American Journal of Psychology*, 114, 329-354.
- Wimber, M., Rutschmann, R.M., Greenlee, M.W., & Bäuml, K.H. (2009). Retrieval from episodic memory: Neural mechanisms of interference resolution. *Journal of Cognitive Neuroscience*, 21, 538-549.
- Wixted, J.T. (2004). The psychology and neuroscience of forgetting. *Annual Review of Psychology*, 55, 235-269.
- Wixted, J. T., Ghadisha, H., & Vera, R. (1997). Recall latency following pure- and mixed-strength lists: A direct test of the relative strength model of free recall. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 23, 523-538.