3-24-2008

A Global Memory Model of Intentional Forgetting

Melissa Lehman
University of South Florida

Scholar Commons Citation

Lehman, Melissa, "A Global Memory Model of Intentional Forgetting" (2008). *Graduate Theses and Dissertations.* https://scholarcommons.usf.edu/etd/357

A Global Memory Model of Intentional Forgetting

by

Melissa Lehman

A thesis submitted in partial fulfillment
of the requirements for the degree of
Master of Arts
Department of Psychology
College of Arts and Sciences
University of South Florida

Major Professor: Kenneth Malmberg, Ph.D.
Douglas Nelson, Ph.D.
Mark Goldman, Ph.D.
Jon Rottenberg, Ph.D.

Date of Approval:
March 24, 2008

Keywords: directed, remember, forget, context, REM

Table of Contents

List of Figures

List of Tables

A Global Memory Model of Intentional Forgetting

Melissa Lehman

ABSTRACT

Intentional forgetting is a phenomenon that has been studied by memory researchers since 1968 (Bjork, LaBerge, & Legrand, 1968), however a formal model to explain directed forgetting has not yet been developed. In this paper, I will review the literature on directed forgetting and discuss the results six experiments used assess directed forgetting in highly controlled manner. The striking findings are a.) that directed forgetting phenomena are observed for both free recall and recognition memory when the list method is utilized, b.) that almost the entire effect in free recall is the result of the ability to initially recall the item from the first serial position, and c.) that the costs and benefits are separately affected by an increase in the retention interval. After extensive model analyses, no simple rehearsal or context based model was identified that can handle the full data set. Here I describe a Retrieving Effectively from Memory model (REM; Shiffrin & Steyvers, 1997) that does account for the full range of findings by blurring the traditional distinctions between these classical approached to directed forgetting phenomena.

Chapter 1: Introduction

Forgetting is one of the most frustrating aspects of daily experience. Sometimes we forget information that we would like to remember and other times we cannot help but to remember information that we would rather forget. The question addressed by my research has been on the latter source of daily frustration. Specifically, I am interested in the cognitive processes involved in *intentional forgetting*.

Forgetting is a hallmark of human memory; it occurs as the result of unconscious, automatic memory processes. However, research suggests that forgetting can also be the result of conscious attempts to control the accessibility of information stored in memory. This form of forgetting is studied in the laboratory using directed forgetting procedures, whereby participants are instructed to forget some material after studying it (Bjork, LaBerge, & Legrand, 1968; MacLeod, 1998 for a review). Memory is then tested for both the to-be-remembered and to-be-forgotten material.

In a *free recall* task, for instance, participants are asked to generate as many items as possible in any order. What is typically found is that to-be-remembered words are remembered better than to-be-forgotten words (Bjork, 1970). On the other hand, other measures of memory show no effect of intentional forgetting; at least this is what is claimed in the literature (cf. Elmes, Adams, & Roediger, 1970). If I take these findings as a given, even if temporarily, they suggest that forgetting is at times under the control of the participant, although not completely so.

There are two methods commonly used to investigate intentional forgetting. In the *item method*, words are presented one at a time, with a cue to remember or to forget each word. In the *list method,* two lists are usually studied and participants are told that they will need to remember both lists (the "remember" condition), or they are told after the first list that they will not need to remember that list because they will not be tested on it later – and that they should try to remember only the upcoming list (the "forget" condition). For example, the experimenter might tell the participant that the first list was only for practice, in order to orient him to the task, and that he would not be tested on that list. Of course, contrary to the instructions, memory is tested for *both* lists, and two effects are found: Participants in the forget condition remember fewer words from the to-be-forgotten list and more words from the to-be-remembered list than participants in the remember condition. These effects are referred to as the "costs" and "benefits" of directed forgetting, respectively. Though I will discuss theories that attempt to explain the findings from both methods, the focus of this paper is on the list method.

Bjork et al. (1968) devised directed forgetting as a method to eliminate proactive interference (PI). In initial experiments, participants were either placed in a remember condition, which had two lists to be remembered; a forget condition, which had one list to be forgotten and another to be remembered; and a no-PI condition, which had only one list. Only the last list from the remember and forget conditions was tested, and this list was compared to the no-PI condition; meaning only the benefits of directed forgetting were examined – and these were observed. Block (1971) used similar procedures to show that directed forgetting eliminates proactive interference, but not retroactive interference; instruction to forget the second list did not improve memory for the first list (again costs were not examined). Using slightly different procedures, Epstein (1969) found that when participants were told after learning a list that they would have to recall only the second half of a list, they performed better on the second half than participants who were told that they would have to recall the second half, then the first half. Eventually, researchers focused on the status of the to-be-forgotten material and began testing memory for the to-be forgotten information. They found that, in addition to benefits, there are costs associated with the instruction to forget (Bjork, 1970).

In a typical *recognition* test, participants are presented with words, some of which were shown on the study list (targets), and some of which were not shown (foils); the task is to decide whether these words were previously studied. In contrast to the findings for free recall, the consensus in the directed forgetting literature is that the costs and benefits associated with the list method are not revealed on a recognition task (see MacLeod, 1998 for a review). For instance, recognition tests often show no effect of the forget instruction (Elmes, et al., 1970; Block 1971; Geiselman, Bjork, & Fishman, 1983; Basden, Basden, & Gargano, 1993), and other times the results have been inconsistent (Sahakyan and Delaney, 2005).

While there has been over 40 years of research and dozens of experiments investigating intentional forgetting under a wide variety of encoding and testing conditions, there is no consensus on how intentional forgetting occurs. In fact, there is no coherent explanation of the effects of even small variations in methodologies. Its not surprising, therefore, that the current literature is unorganized and there is no clear direction for research to proceed. The goal of this research is to establish a theoretically relevant set of benchmark empirical findings concerning the effect of instructions to forget and develop a global model embedded in a rich theoretical framework to explain how directed forgetting occurs and to interpret my findings.

Classical Directed Forgetting Hypotheses

Before I discuss the hypotheses regarding the way that directed forgetting occurs, I must discuss the possibility that participants are not actually forgetting the words, but they are failing to output the words due to demand characteristics of the task. Researchers have proposed that participants are not recalling the to-be-forgotten words because they believe they should have forgotten them after the instruction. In order to eliminate this possibility, some researchers have used money as motivation and still found costs and benefits. Woodward and Bjork (1971) first had participants recall the to-be-remembered words, awarding them $.01 for each one they recalled, and penalizing them $.01 for each to-be-forgotten word that they recalled. Participants recalled about half of the to-be-remembered words, and less than 2% of to-be-forgotten words. After multiple lists were learned, participants were asked to recall both to-be-remembered and to-be-forgotten words, and they were awarded $.01 for both types of words. Woodward and Bjork found performance similar to that in the first task; even when motivated by money, participants were unable to recall the to-be-forgotten words. MacLeod (1999) replicated these findings; when offered $.50 for each of the to-be-forgotten words recalled, he saw almost no improvement in memory for these words. Given that demand characteristics cannot explain the effect of forgetting, I will next discuss hypotheses involving the way that participants are able to forget.

Many of the earliest hypotheses were developed by R. A. Bjork and his colleagues (1968). Even though their initial research considered only the benefits of directed forgetting, all of their hypotheses have also been applied to explain the costs. For instance, the *erasure hypothesis* states that participants effectively erase items from memory when given the forget instruction. As a result, the to-be-forgotten words are no longer present in memory, and compared to participants who have not had this material erased, performance is worse, hence the costs of directed forgetting. Additionally, because these words have been erased from memory, they are no longer available to create proactive interference, and memory on the following list is better – hence the benefits.

Bjork et al. (1968) suggested that the erasure explanation, while plausible, was not likely. Bjork (1970) tested to-be-forgotten items and found that while memory for these items was worse than for to-be-remembered items, some were still remembered – unlikely if erasure is possible. Other research suggests that memories do not exist in an all-or-nothing state, which is implied by the erasure hypothesis (Atkinson, Bower, & Crothers, 1965). Moreover, the erasure hypothesis cannot predict which items would be erased from memory and which items would not be. Lastly, the hypothesis cannot simultaneously predict why erasure would sometimes occur for free recall but never for recognition. I will not consider the erasure hypothesis any further for present purposes.

The Differential Rehearsal Hypothesis

A more viable hypothesis is the *differential rehearsal hypothesis* (Bjork et al., 1968), which says that participants in the forget condition stop rehearsing words from the to-be-forgotten list after the forget instruction and devote all further rehearsals to the following list, but this does not occur for participants in the remember condition. Because the words on the to-be-remembered list receive comparatively more rehearsals after instruction to forget, they are encoded better. Because words from the to-be-forgotten list are encoded less well, they are assumed to decay at a greater rate, leading to the costs of directed forgetting. As with the erasure explanation, these words are no longer available, and thus do not create proactive interference, leading to the benefits.

The differential rehearsal hypothesis, at least this form, is unlikely to provide a complete explanation of directed forgetting for several reasons. First, items encoded at varying levels of initial strength are nevertheless forgotten at similar rates (e.g., Ebbesen & Wixted, 1991; Slamecka & McElree, 1983). Second, directed forgetting is commonly observed even though the conditions of the experiment are specifically designed to prevent rehearsals. For instance, Bjork et al. (1968) and Block (1970) had participants perform various shadowing tasks during study. On the assumption that participants could not simultaneously selectively rehearse first-list items and attend to the shadowing task, it is unclear why they observed the effect of instructions to forget. Geiselman, Bjork, and Fishman (1983) investigated intentional versus incidental learning. They assumed that differential rehearsal would not be engaged when memory was incidentally tested and found the costs and benefits of directed forgetting for both intentionally and incidentally learned words. They argued that while selective rehearsal could explain the directed forgetting effect on the intentionally learned items, it does not explain the effect on the incidentally learned items.

It is important to mention that differential rehearsal might help explain the directed forgetting effect for the item method. Whereas for the list method, participants are required to rehearse the words from the list because they do not find out until after the initial study list that they will not need to remember them, in the item method, participants can decide during learning whether they will rehearse a word – because the cue is given after each word is presented. MacLeod (1975) found that when using the item method, the directed forgetting effect persists over a one-week delay, evidence he used to support the rehearsal explanation. I will address the issue of delay later in this paper.

The erasure and rehearsal explanations are similar in that they both posit that the mechanism behind directed forgetting is a failure to completely encode to-be forgotten information. Alternatively, the *set differentiation* hypothesis assumes that the forget instruction has its effect after the information has been encoded. Bjork et al (1968) proposed that participants respond to the forget instruction by differentially coding to-be-remembered and to-be-forgotten information, in a way that reduces interference between the two. The set differentiation hypothesis states that participants effectively group the to-be-remembered and to-be-forgotten items separately (Bjork, 1970). When given the forget instruction, they differentiate the to-be-forgotten words by putting them in a separate group from the to-be-remembered words that follow the forget instruction. The set differentiation hypothesis differs from erasure and rehearsal hypotheses because while the set differentiation hypothesis assumes that items are differentiated during encoding, the items *are* completely encoded, but into different sets, and during retrieval only the to-be-remembered set is searched.

It is important to note that none of these hypotheses are necessarily mutually exclusive. For instance, Bjork (1970) assumed a combination of the set differentiation and rehearsal explanations; participants group the to-be-remembered and to-be-forgotten items separately, and then devote all rehearsal to the to-be-remembered group. Even so, the additional complexity associated with a combination of these hypotheses in their original form cannot account for the differential effects of intentional forgetting on free recall and recognition.

Geiselman, Bjork, and Fishman (1983) suggested that *retrieval inhibition* was the process responsible for the directed forgetting effect. According to the retrieval inhibition hypothesis, participants effectively group the to-be-forgotten and to-be-remembered material separately, and then inhibit the to-be-forgotten set. While these items are still present in memory, they are inhibited during retrieval, thus leading to the costs of directed forgetting. Because they are inhibited, these items do not create proactive interference, leading to the benefits. Elmes, Adams, and Roediger (1970) suggested that the to-be-forgotten information is "suppressed" while to-be-remembered information is "selected". Again, however, the retrieval inhibition hypothesis cannot explain the effects of intentional forgetting on free recall and recognition. Moreover, neither explanation tells us much about how this process works; when participants differentiate the words into two separate sets, how are they able to "inhibit" one set and "activate" the other?

The Contextual Differentiation (CD) Hypothesis

Sahakyan and Kelley (2002) proposed a different explanation of directed forgetting based on Bjork's (1970) set-differentiation framework. According to their hypothesis, study involves the storage of information representing the studied items (i.e., item information) and the context in which the items occur (i.e., context information). Each list is associated with an overlapping but not completely similar set of contextual elements. Sahakyan and Kelley hypothesized that after receiving a forget instruction, participants engage in a mental context change. This causes the overlap between the contextual elements associated with list 1 and list 2 to decrease. As a result, the contexts of list 1 and list 2 are more contextually differentiated in memory than they would be without the forget instruction. In addition, the list 1 context is less similar to the context at test after an instruction to forget due to the change in mental context that occurred. When recalling list 2, there is therefore less interference from the list 1 traces. Hence, this is the source of the benefits of the instruction to forget. The costs are the result of the relative inaccessibility of an effective context cue for list 1 traces, again due to the change in mental context that occurred between the list presentations.

The context model of directed forgetting has its basis in a large literature on the effects of context change on memory performance. Godden and Baddeley (1975) used an environmental context change to alter memory performance on a recall test. Participants learned a list of words either on land or under water, and then were tested in either the same or a different context. Memory performance was impaired when test context was different from study context. Similar impairments have been found in studies that used a mood-context change or a state-context change. Macht, Spear, and Levis (1977) showed that when participants were tested on words while in a different mood state than their mood state during study (i.e. anxious vs. calm), performance was worse than when mood during study matched mood during test. Similarly, researchers have found that participants who studied while under the influence of alcohol or marijuana performed better when they were under the influence during test than when they were not (Goodwin, Powell, Bremmer, Hoine & Stern, 1969; Eich, Weingartner, Stillmin & Gillin, 1975, respectively).

Based on the context-change literature, Sahakyan and Kelley (2002) hypothesized that participants given the forget instruction undergo a mental context change. In pilot work, they asked participants to retroactively report on their strategies when given the forget instruction. They report that participants often claimed to "think about something else" – a method of changing internal context. In order to test the context-change hypothesis, they designed an experiment where half of participants participated in standard directed forgetting conditions, and half participated in a context-change condition. In the context change condition, participants were given either the remember or forget instruction, followed by an instruction to change mental context. Participants were instructed to imagine that they were invisible, and to think about what they would do if they would suffer no consequences for their actions. There were no differences in performance between the remember-plus-context-change (RCC) and forget-plus-context-change conditions, so I will discuss data from only the RCC condition.

Participants in the RCC condition performed almost identically to participants in the forget condition – showing both costs and benefits of the context change, compared to participants in the remember condition. Sahakyan and Kelley took this as support for the context change explanation of directed forgetting.

Context Reinstatement

Smith (1979) showed that while participants who change rooms between study and test show impairment compared to participants who are tested in the same room that they studied in, participants tested in a different room who mentally reinstate the environmental context of the study room do not show this impairment. Participants studied a list of words in a one room, and some of participants were switched to a different room for test. Half of participants who were tested in a new room were given the instructions to write down ten things they remember about the room that they studied in, and to remember their thoughts, feelings, and the sensations that they experienced in the study room. Participants in this reinstatement condition performed as well on the task as participants who were tested in the original study room.

In a second context change experiment, Sahakyan and Kelley examined the effect of context reinstatement on directed forgetting. At the beginning of the experiment, they played music from the movie *Star Wars* in order to create a distinct context. They again used standard remember and forget conditions, along with the remember plus context change (RCC) condition. After studying the second list, half of participants participated in a context reinstatement procedure – they were instructed to imagine what they were doing immediately before the experiment, and describe their thoughts and feelings as they entered the room, along with what they remember noticing about the room or the experiment. After receiving the context reinstatement instructions, participants in the forget and RCC groups showed significantly reduced costs and benefits compared to the groups that did not receive the reinstatement. While the costs and benefits were not completely eliminated, they were certainly reduced – and if there were a perfect way to mentally reinstate list 1 context, perhaps they could be completely eliminated. These findings revealed not only that context reinstatement has similar effects in the directed forgetting paradigm to those in the environmental context change paradigm, but also that the to-be-forgotten words were still present in memory (further evidence against the erasure and rehearsal explanations, which suggest that information is not completely encoded in the first place).

Much of the data from previous directed forgetting studies can be better explained by the context change explanation. The Geiselman, Bjork, and Fishman (1983) intentionality manipulation study was explained by retrieval inhibition. Perhaps a more clear explanation would be that participants "inhibited" the to-be-forgotten information by utilizing an internal context change. After they were given the forget instructions, participants switched their mental context (possibly by thinking about something else). As a result, the context of list 1 was dissimilar to the test context – and the incidentally learned items from list 1 were part of this dissimilar context.

Additionally, much of the data regarding intrusion errors supports the context change hypothesis. For example, Bjork (1970) looked at intrusion rates in a cued recall task and found that intrusions were very rarely to-be-forgotten items; they were almost always to-be-remembered items. Bjork manipulated the number of to-be-remembered and to-be-forgotten pairs (from 1 to 5 and 0 to 3, respectively). As the number of to-be-remembered pairs increased, number of intrusions from these pairs increased; however, an increase in the number of to-be-forgotten pairs did not increase the number of intrusions from those pairs. The forget instruction leads to differentiation between the contexts of the to-be-forgotten and to-be-remembered lists – so, even though participants are able to remember some of the to-be-forgotten words, these words barely intrude because participants are better able to differentiate between the to-be-remembered and to-be-forgotten information. Because the to-be-forgotten pairs occur in a different context, they will not intrude on the to-be-remembered pairs, regardless of the number that are presented in this different context. Other to-be-remembered pairs, however, occur in the same context as the tested item, thus the amount of interference they create will increase as the number of pairs increases.

Recall and Recognition

Like all prior hypotheses, the context-change hypothesis cannot readily explain why intentional forgetting affects free recall but does not affect recognition. The goal here is to provide a simple account of intentional forgetting that explains the effect of directed forgetting for both tasks. I began this research by noting that extant literature consistently reports poorly designed experiments. In the following sections, I describe an improved design and report new findings that test several more specific assumptions about how a change in mental context might allow for intentional forgetting.

Importantly, the new model predicts that intentional forgetting should affect recognition as well as free recall. Given that I found that extant free recall experiments were poorly designed, I suspected the same would be true of the prior recognition experiments. Indeed, this was the case and the problems were compounded by poor measurement instruments. As a result, I have conducted several recognition memory experiments using the improved design and the proper measures.

Chapter 2: Experiments

Experiment 1 – Free Recall

According to models of free recall (Malmberg & Shiffrin, 2005; Raaijmakers & Shiffrin, 1980), context plays a major role during retrieval. When asked to recall freely a list of recently studied items, for instance, the retrieval cue used to probe memory consists of mentally reinstated contextual elements. The effectiveness of these cues is a positive function of the similarity between the context in the retrieval cue and contents of memory. Most models of contextual dynamics assume that the contextual elements available to be encoded change over time (Estes, 1955; Mensink & Raaijmakers, 1989; Howard & Kahana, 2002). Thus, these models make the straightforward prediction that more recent events should be better remembered than less recent events. Long-term recency effects are commonly found in memory literature (e.g., Ebbinghaus, 1885).

A recency effect is therefore a critical prediction of any model; list 2 (L_2) should be remembered better than list 1 (L_1) all thing being equal. A review of directed forgetting literature shows, however, the *opposite* is almost always the case. Thus, it is possible that the CD model can be rejected based on prior findings. However, all things might not be equal. The lack of a recency effect in prior experiments might be due to some experimental confounds that aid encoding of L_1 while impairing encoding of L_2. For instance, the list method usually utilizes only two lists; L_2 experiences proactive interference from L_1, but L_1 has no such list before it to create interference. In addition, L_2 is usually followed by a distractor task, which prevents rehearsal of L_2 items, whereas

L_1 is not (Brown, 1958; Glanzer & Cunitz, 1966; Peterson & Peterson, 1958; but see Bjork & Whitten, 1974). This might encourage participants to continue rehearsing words from L_1 while they are learning L_2, to the detriment of L_2 words. To control for these confounds, I utilized a three-list design (cf. Jang & Huber, 2008; Sahakyan, 2004), with a distractor task after each list. The forget instruction comes after L_2 and participants are only tested on lists 2 and 3. This ensures that each list is preceded by another list and followed by a distractor task.

A second prediction concerns intrusion errors. Because context at test is more similar to the context of L_3 than to the context of L_2, the number of intrusions from L_3 while trying to recall L_2 should be greater than the number of intrusions from L_2 when trying to recall L_3. Another prediction concerning intrusion errors is that intrusion rates will be lower in the forget condition than in the remember condition. Because the context change that occurs with the forget instruction makes the lists more distinct, participants should be better able to determine whether a retrieved word came from the wrong list and thus fail to output that word. This prediction is consistent with Bjork's (1970) findings that intrusions almost always come from to-be-remembered rather than to-be-forgotten material. Finally, the number of intrusions of an L_1 item will be greater when trying to recall from L_2 than when trying to recall from L_3. The context used to probe memory for L_2 will be more similar than L_3 will be to L_1. Additionally, the model predicts fewer L_1 intrusions after the forget instruction because the context at test will be less similar to the context of L_1, regardless of whether L_2 or L_3 is to be recalled (cf. Sahakyan, 2004). However, intrusions errors are usually quite rare. Thus, I might not be able to gather meaningful data to test these predictions using a free recall procedure. Intrusion errors are much more common when a recognition procedure is used. I will return to address this methodological issue when I discuss the recognition experiments that are to follow.

These predictions can be contrasted with those of the differential rehearsal hypothesis. Usually words studied at the beginning of a list are remembered better than words studied at later serial positions when memory is tested after some filled delay (i.e., a *primacy effect*). According to many models, the primacy effect is observed because early-list items are rehearsed longer than later-list items (Atkinson & Shiffrin, 1968). Additionally, words at the end of a list will be rehearsed in short-term memory until test time. According to the same models, this produces the recency effect. As I mentioned earlier, however, a distractor task eliminates the recency effect, by preventing participants from rehearsing words from the end of the list (Glanzer & Kunitz, 1966). If the differential rehearsal explanation for directed forgetting is correct (Bjork et al., 1968), participants in the remember condition covertly rehearse L_2 items while beginning to study L_3, but this should not occur in the forget condition of the experiment because presumably there is no reason to continue studying the L_2 items. If so, I should observe no, or perhaps an attenuated, primacy effect on L_3 in the remember condition. Moreover, if participants in the remember condition are continuing to rehearse words from the end of L_2 while they are learning L_3, then I should see a recency effect for L_2.

Method

Participants. Participants were 180 undergraduate psychology students at the University of South Florida who participated in exchange for course credit. Data for twelve participants were not used because they were unable to recall any words from any lists leaving 168 participants (42 per condition).

Materials. Experiments were all run using Authorware software, which allows for presentation of visual information and input of user responses. The entire experiment was completed on a computer in an individual participant room. For each participant, 48 words were randomly chosen from the Francis and Kucera (1982) norms and divided into three lists of 16 words.

Procedure. At the beginning of the experiment, participants were shown an information page that told them about the supposed purpose of the study. They were told that the experimenters wanted to see how well people could not only remember information but also remember where that information came from. Participants were informed that they would see three lists of words, and that they would be tested on only one of the lists, but they would not be told which list until later in the experiment, so they needed to remember all of the lists. The instructions were as follows:

> *At the beginning of this experiment, you will study three lists of words. The words will appear on the screen one at a time for a few seconds each.*
>
> *Your task is to remember these words for a later memory test. Importantly, I will only ask you to remember the words from one of the lists, which will be chosen randomly, but you will not be told which list until later in the experiment.*
>
> *In between each list there will be short math task. This is involves adding digits in your head and entering the total into the computer. Once you have done so, the next list of words will be presented.*

Once they understood the instructions, they continued on to begin the study lists. Participants were given a warning before each list that the study list was about to begin.

Lists were shown one word at a time, with each word appearing on the screen in black on a white background for 8 seconds. The lists consisted of 16 words. After each list, participants participated in a math distractor task, where they completed two-digit addition problems. The distractor task lasted 30 seconds and participants were instructed to complete as many problems as they could in this amount of time.

Participants in the remember condition were shown each list and distractor task followed by the test. In the forget condition, participants were shown the first two lists and distractor tasks. They were then shown the forget instruction, followed by study of the third list and a third distractor task. The forget instruction was as follows:

> *Next you are going to receive the third study list. This is the list that you will be asked to recall, so you do not need to worry about the first two lists.*

After all three study lists (and distractor tasks), participants were given a free recall test lasting 90 seconds. They were told to enter onto the screen all of the words that they could remember from the specified list. Half of participants in each condition were tested on L_2 and half were tested on L_3. Participants from the forget condition who were tested on L_2 (the "forget") list were told that I want them to recall from this list even though I had previously told them that they won't need to remember it. After being tested on the specified list, participants were tested on the other list (either list 2 or list 3), however this data was only used to determine whether any participants failed to recall any words from either list – in which case their data was thrown away.

Results

Correct Recall. The statistical analyses are confined to the data obtained from L_2 and L_3. The results of a two-way ANOVA show that for correct recall, there was a main effect of list, $F(1,164) = 68.84$, $MSE = .021$, $p < .001$; for both remember and forget conditions, probability of recall was significantly greater for L_3. This confirms the prediction that more recent lists will be better remembered than less recent lists. There was no main effect of instruction, but there was a significant List x Instruction interaction, $F(1,164) = 19.66$, $p < .001$. As shown in *Figure 1*, recall was better for the remember condition than the forget condition on L_2, but the opposite is true on L_3. According to planned comparisons, all results shown here are significant to a .05 criterion.

Figure 1. Probability of correct recall and intrusion errors for free recall (Experiment 1).

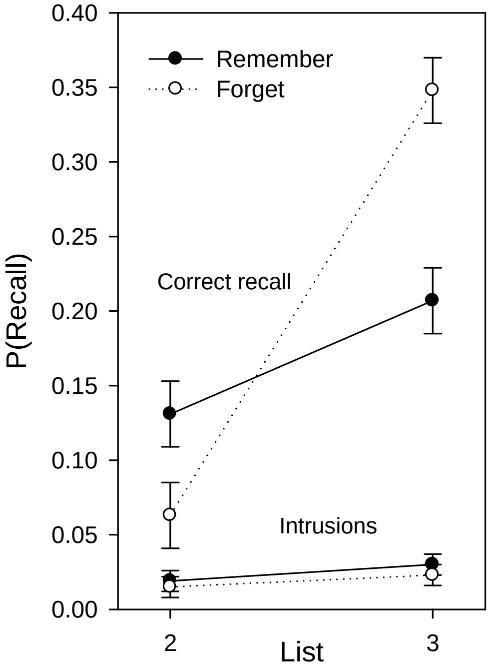

Note: The intrusions in this graph refer to intrusions that came from either list 2 or list 3. When recalling from list 2, any list 3 item that was output is referred to as an intrusion and vice versa.

Intrusions. As expected, intrusion rates were very low. These are shown in *Figure* 1. For intrusions from L_2 when recalling L_3 and vice versa, there were no significant results; however there were some interesting trends. Participants were more likely to have intrusions from L_3 while being tested on L_2 than they were to have intrusions from L_2 while being tested on L_3. Further, the probability of either type of intrusion is lower for participants in the forget condition than in the remember condition. For intrusions that came from L_1, there is a significant main effect of List, $F(1,164) = 17.83$, $MSE = .002$, $p < .001$; participants were more likely to have intrusions from L_1 while they were recalling L_2 than when they were recalling L_3 for both remember and forget conditions, and again intrusion rates were lower for participants in the forget condition. These are shown in *Table 1*.

Table 1.

List 1 intrusions

	Remember		Forget	
	L_2	L_3	L_2	L_3
Free Recall	0.051	0.016	0.038	0.01
Delay	0.0938	0.0327	0.0682	0.0556

False-alarm rates for unstudied foils

	Remember		Forget	
	L_2	L_3	L_2	L_3
Exclusion 8s	0.112	0.083	0.143	0.073
Exclusion 4s	0.1584	0.0778	0.2151	0.0611

False-alarm rates for inclusion

	Remember	Forget
Inclusion 8s	0.0792	0.0771
Inclusion 4s	0.1003	0.0683

Note. In inclusion, L_2 and L_3 were tested together, thus there are only false-alarm rates for remember and forget conditions.

Thus, the model correctly predicts that intrusions from L_3 while being tested on L_2 are greater than intrusions from L_2 while being tested on L_3, intrusion rates are lower in the forget condition than in the remember condition, and intrusions from L_1 are more likely when participants are recalling from L_2 than when recalling from L_3.

Serial Position. In addition to looking at correct recall and intrusion errors, I examined serial position data. This allowed us to explore in detail the contribution of rehearsal to the directed forgetting effect. The left panel of *Figure* 2 shows the serial position functions obtained in my experiment. There is a primacy effect for L_3 in both the forget and remember conditions, and it is smaller in the remember condition than in the forget condition. There was a significant main effect of Serial Position, $F(15,164) = 5.59$, $MSE = .125$, $p < .001$. The List x Instruction x Serial Position interaction was not reliable. The smaller primacy effect for L_3 in the remember condition suggests that participants might not have given as many rehearsals to words at the beginning of L_3 as in the forget condition because they are still rehearsing words from L_2.

I also examined another aspect of serial position to further understand the rehearsal component of the directed forgetting effect; the first item output during recall. There was a significant main effect of Serial Position, $F(15,161) = 10.04$, $MSE = .057$,$p < .001$, and this was moderated by a significant List x Instruction x Serial Position interaction, $F(15, 161) = 3.29$, $p < .001$ for first item recalled at test. The right panel of *Figure* 2 shows that the first item recalled from L_2 was most likely to be the first word on the list only in the remember condition, whereas on L_3, the first item recalled was more likely to be the first word on the list for the forget condition than the remember condition.

This indicates that the instruction to forget had a large impact on which word would be recalled first, and L_2 and L_3 were impacted in opposite directions.

Figure 2. Serial position data for free recall (Experiment 1).

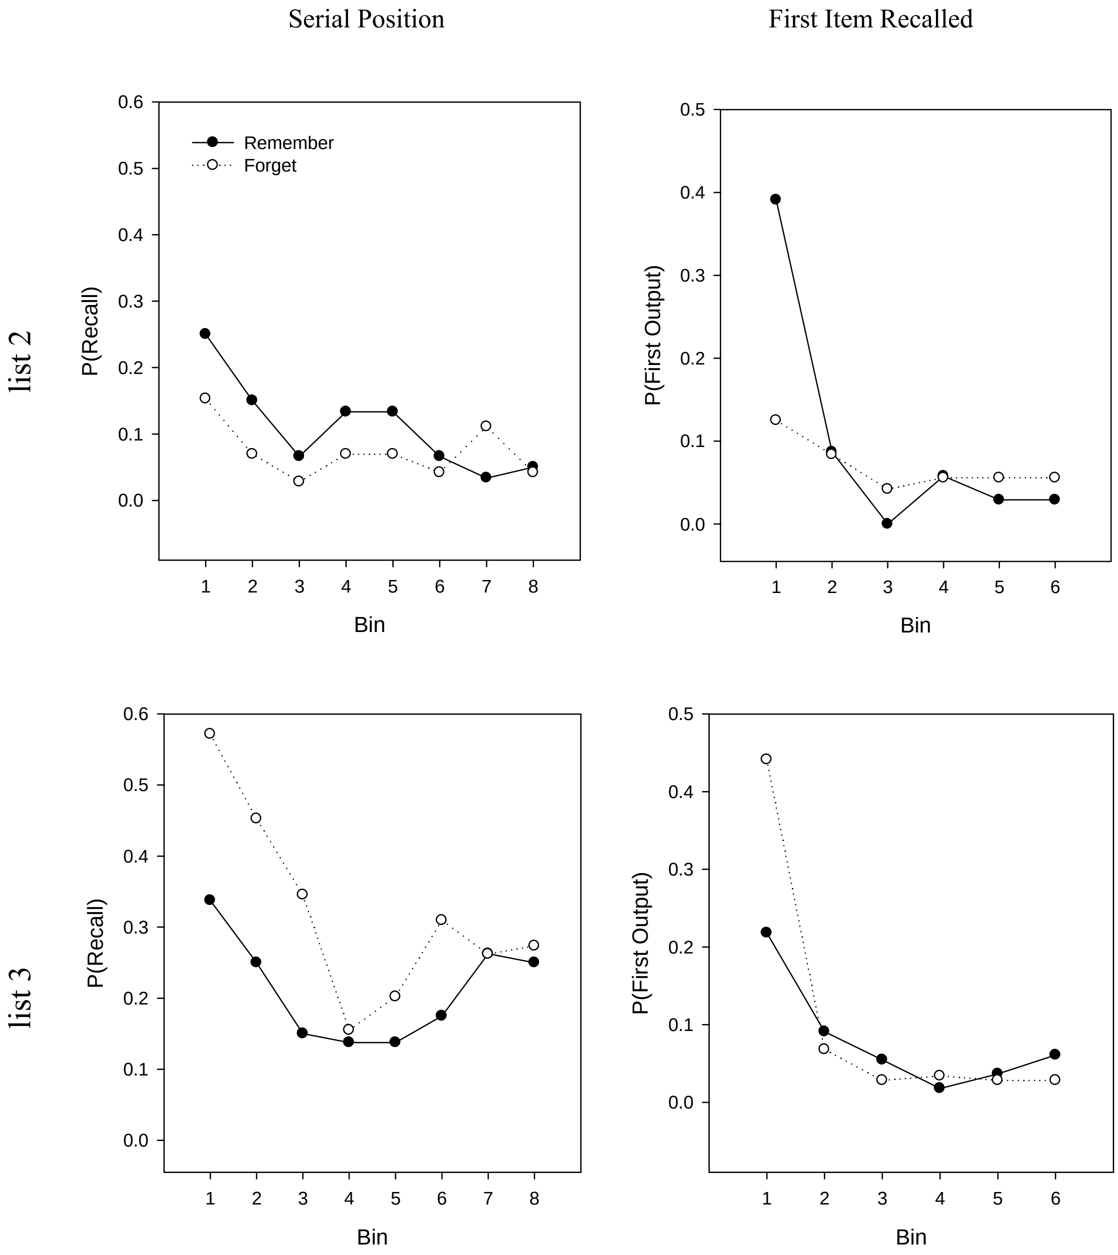

Note: For the sake of clarity, the 16 item list was compiled into bins. For serial position, each bin spanned two serial positions. For instance, bin *n* contains the data from serial positions *2n-1* and *2n.* For first item output, bin 1 represents the first item on the list (since this is where differences are seen) and all other serial positions are grouped by three.

In addition to looking at serial position and first item output patterns to evaluate the contribution of rehearsal, I also examined the conditional response probability for each condition. Conditional response probability (CRP; Kahana, 1996) refers to the probability of recalling an item from a given serial position after successful recall of an item from a nearby or distant serial position. The distance in serial position from one item recalled to that of the next item recalled is referred to as *lag*, and lag can be forward or backward. For example, if a participant recalls an item from the 5th serial position, and next recalls an item from the 6th serial position, this would be represented by a lag of +1. If the participant next recalls an item from the 2nd serial position, this would be represented by a lag of -4. Conditional response probabilities allow us to see the degree to which participants are recalling successive items on a list successively – presumably this indicates that they are using previous items as cues with which to probe memory. As shown in *Figure* 3, CRP curves differ slightly between conditions, $F(1,87) = 2.107$, $MSE = .092$,$p < .001$. On L_2, participants in the forget condition are more likely to move in the backward direction than participants in the remember condition.

Figure 3. Conditional response probabilities for free recall (Experiment 1).

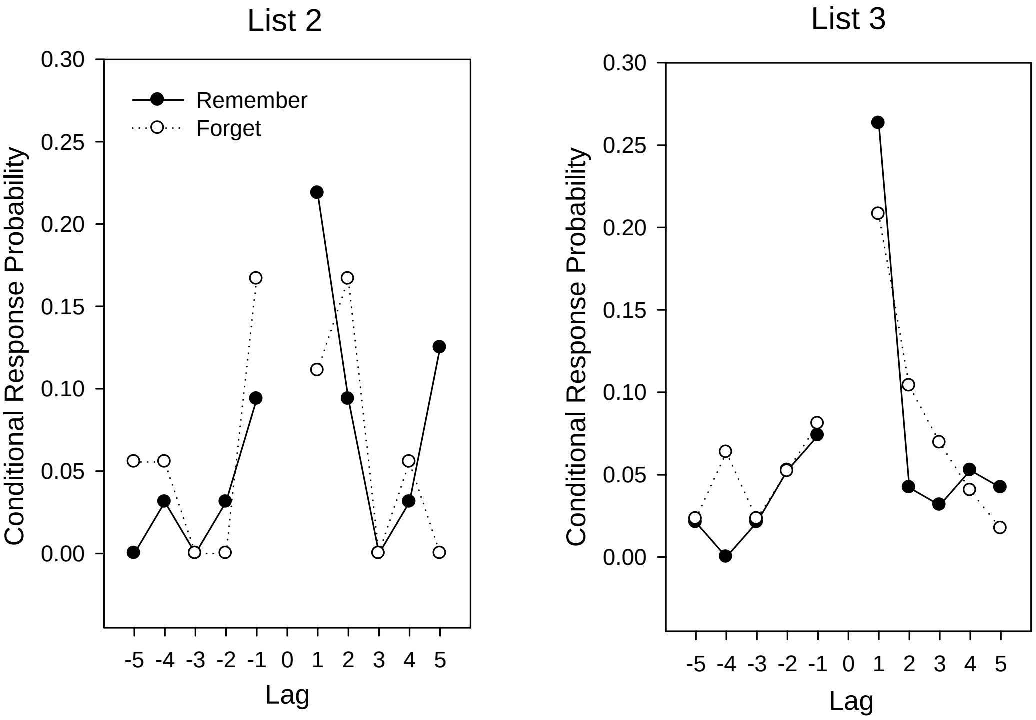

Note. Conditional response probabilities refer to the probability of recalling an item a given distance away from the previous item in serial position. Successive recall of nearby items is more likely than recall of distant items, and forward movement is more likely than backward.

Discussion

I simultaneously observed the cost and benefits associated with directed forgetting and a recency effect. Additionally, while intrusion rates were extremely low, the patterns of intrusions seen in this experiment are consistent with a context differentiation hypothesis.

Serial position effects, on the other hand, suggest that a context differentiation hypothesis is not enough to explain the difference between conditions. A difference in the amount of primacy on L_3 suggest that participants are differentially rehearsing the material in remember and forget conditions. While the primacy effects were consistent with the differential rehearsal hypothesis, however , other aspects of the serial position curves were not. For instance, if the differential rehearsal hypothesis is correct, then the costs and benefits should have only been observed for the initial serial positions. However, the cost and benefits were observed for almost the entire length of serial position curves. In addition, I failed to find any recency effect on L_2 for either the remember or forget conditions. This suggests that participants in neither the forget condition nor in the remember condition were rehearsing words from the end of L_2 while they are learning L_3.

These findings suggested, however, that it was possible that while participants were rehearsing words from L_2 while studying L_3, these words did not come from the end of L_2. That is, it is possible that the distractor task was effective in preventing rehearsal during the interval between the lists, but at the beginning of list L_3 other L_2 items were retrieved from memory and rehearsed. To evaluate this new hypothesis, I sought to determine which items from L_2 would be most likely to be retrieved after the distractor task and co-rehearsed at the beginning of L_3. It is likely that the word that is getting the most rehearsals during study will be the first word output during recall (Rundus, 1971). The first-item output analysis revealed a different pattern between remember and forget conditions, and suggested that although participants in the remember condition probably were not rehearsing words from the end of L_2 while they were learning L_3, they may have been rehearsing words from the beginning of L_2. Additionally, when recalling L_3, they were not very likely to first output the first item from the list – suggesting that this item is not getting as many rehearsals as it would if L_2 was not being rehearsed during study of L_3.

The serial position data supports a differential rehearsal explanation. However, the observed serial position data are also consistent with the context model if the context cues at test are more strongly associated with the items at the beginning of each list. This would produce an advantage for first items on the list when the first attempt at recall was made. This advantage should diminish for L_3 in the remember condition due to the additional competition from L_2 items. In contrast, the first recall advantage for L_2 items should be greater following the instruction to remember, since there is less contextual drift and thus the appropriate context cues are more readily available.

The CRP curves in this experiment are consistent with those typically found in free recall experiments – participants are most likely to successively recall from nearby serial positions, and recall is more likely to move in a forward than a backward direction. These are referred to as contiguity effects, and according to some models they are due to the fact that items from nearby serial positions are associated with a more similar set of contextual elements than those from distant serial positions (Howard & Kahana, 2002).

The CRP findings are also consistent with my serial position data. Participants in the remember condition are more likely than participants in the forget condition to recall the item in the first serial position on L_2. The CRP analysis showed that participants in the forget condition were more likely to move backward on L_2 than participants in the remember condition. Since participants in the remember condition are more likely to recall the item first item on L_2, they can only move forward, thus will have a lower number of backward recalls.

The CRP functions and the relative lack of recency for L_2 in the forget condition is a challenge for the differential rehearsal model to explain. At the same time however, the first-item output findings and the lack of primacy for L_3 in the remember condition suggest that differential rehearsal may be playing a part in the directed forgetting effect. Of course, as a REM is a descendent of Atkinson and Shiffrin (1968) modal model, assumptions that produce differential effects of rehearsal an important component of the model. The REM based model I have developed assumes that rehearsal affects the storage of item information and context information. I will discuss this model in detail after I present the results from the remaining experiments.

Experiments 2 and 3 – Recognition: Exclusion versus Inclusion

It is widely believed that directed forgetting does not affect recognition memory (MacLeod, 1998), but I am skeptical for several reasons. Both recognition and free recall are episodic memory tasks, and there are very few variables that affect one but not the other. In addition, the recognition memory experiments that have been reported often used measures and procedures that are less then optimal, and most of these studies did not report all of their recognition data. For example, Block (1971) reported *d'*. This tells us that sensitivity was unaffected, but it does not tell us anything about bias, which may have been affected by the forget instruction. Other times small differences in recognition performance were observed, but they were not statistically significant. Note that context effects on recognition can be very small (cf. Murnane, Phelps, & Malmberg, 1999), and therefore it is possible that the statistical tests used in these experiments are underpowered. For instance, it is common to use a sample size of 40+ in each condition of a directed forgetting experiment (cf. Sahakyan and Kelley, 2002), but in recognition experiments, samples have been as small as 16 participants in each condition (cf. Basden, Basden, & Gargano, 1993).

Many of the methodological issues were probably the result of the fact that recognition memory has never been investigated in a systematic fashion. In fact, many (perhaps most) of the experiments that have been reported simply threw in a few recognition memory trials at the end of a free recall experiment. For instance, Basden et al. (1993) compared to-be-forgotten words to to-be-remembered words for participants in the forget condition, but did not have a remember condition to which to compare these words. Often, the recognition test followed a free recall test (eg. Elmes et al., 1970). This is problematic because it could add noise to the data, meaning more variation and less ability to see an effect.

Perhaps most importantly, nobody has developed a cogent model that can predict why free recall and recognition might be differentially affected by directed forgetting. In fact, all models predict that recognition should be affected by directed forgetting, including the model that I have developed. Interestingly, Sahakyan and Kelley (2002) use a failure to find effects of context change on recognition memory (Godden and Baddeley, 1980) as support for the context change hypothesis of directed forgetting. There are, however, problems in the literature on context change and recognition. As in the directed forgetting recognition literature, researchers in the field of context change and recognition often fail to report all of their data. They report that there is no difference in sensitivity (*d'*); however this does not tell us whether there was an effect of the context change. Indeed, context dependent recognition has been consistently observed when the appropriate experiments have been conducted (cf. Murnane et al., 1999). In any case, the context change account of directed forgetting *does* predict an effect in recognition, and data showing that context change does not have an effect on recognition actually disconfirms this account.

Given the shortcomings of the prior literature, it is worth considering the different ways in which recognition memory can be tested, particularly since the list method requires multiple study lists. Under these conditions, recognition experiments can use either an *inclusion* test or an *exclusion* test (Jacoby, 1991; Winograd, 1968). In an inclusion test, the task is simply to say "yes" if a word was studied on any list during the experiment. Hence, the context that differentiates the study lists is not required in order to perform the task. Indeed, it is logical to assume that the context used to probe memory would tend to consist of those context features that the study lists have in common. Compare this to the free recall procedures used in the direct forgetting literature where participants must use a context cue for a particular list, the one specified by the experimenter.

The different context cues required to perform free recall and inclusion tasks might explain the different effects of directed forgetting that have been reported. In any case, an inclusion test is not the optimal test for examining context effects on memory. Interestingly, all of the recognition memory experiments in the directed forgetting literature have used an inclusion procedure. Thus, if I accept a CD account for directed forgetting, then it is reasonable to expect small or perhaps even null effects of directed forgetting on recognition performance. While some may consider an inclusion task to be a purely based on a judgment of familiarity of a test item, there are still contextual elements that will affect memory on this task. As context is used in the retrieval process, a more recent list should be better remembered than earlier lists, because context at test will still be more similar to that of the last list (as it was in the recall experiment). Additionally, because the forget instruction differentiates context at encoding, the context change that it elicits should not differ between recall and recognition tasks. The context change makes the context at test less similar to that of the to-be-forgotten list, increasing the difference between memory for L_2 and L_3 in the forget condition. Note that because the performance of an inclusion task may not rely on contextual elements that differentiate the lists, the effects of the forget instruction may be small. Nevertheless, they should still be observed. Based on the CD model's predictions, I again expect to see costs and benefits, along with recency of L_3.

In an exclusion test, on the other hand, the participant's task is to say "yes" only if a word was studied on a single specified list. In this case, the participant must use context in the retrieval cue that differentiates the study lists in order to accurately perform the exclusion task. Thus, if the context model is accurate, then I would expect to see robust effects of directed forgetting on exclusion task performance. That is, because the exclusion test is a context-based task, I expect to see effects similar to those in free recall, including the costs and benefits of directed forgetting on hit rates and the recency of L_3 in the remember condition. Additionally, false alarm rates should be similar to the intrusion rates for the free recall experiment; there should be more L_3 false alarms when a participant is attempting to recognize from L_2 than there will be L_2 false alarms when a participant is attempting to recognize from L_3, and intrusion rates should be lower in the forget condition. In order to facilitate modeling, the following experiment utilized a design exactly like that used in the prior free recall experiment except that recognition memory is tested using an exclusion procedure.

Methods

Participants, Materials, and Procedure. In the exclusion experiment, participants were 148 (37 in each condition) undergraduate psychology students at the University of South Florida who participated in exchange for course credit. For each participant, 64 words were randomly chosen from the Francis and Kucera (1982) norms to create four lists of 16 words. The study procedure was identical to that of Experiment 1. For the test, participants were told which list they would need to recognize words from (either list 2 or list 3). They were told to respond “yes” if the word shown was on the specified list, and to respond “no” if the word shown was from a different list or if it was a new word. The test list consisted of all of the words from lists 2 and 3, and 16 new words, in a random order.

In the inclusion experiment, participants were 60 (30 in each condition) undergraduate psychology students at the University of South Florida who participated in exchange for course credit. The design, material, and procedure was the same as in the exclusion experiment, except participants were told that if they had seen that word on either list (2 or 3) then they should respond by clicking “yes” and if the word was a new word they were to respond by clicking “no.”

Figure 4. Exclusion recognition performance (Experiment 2).

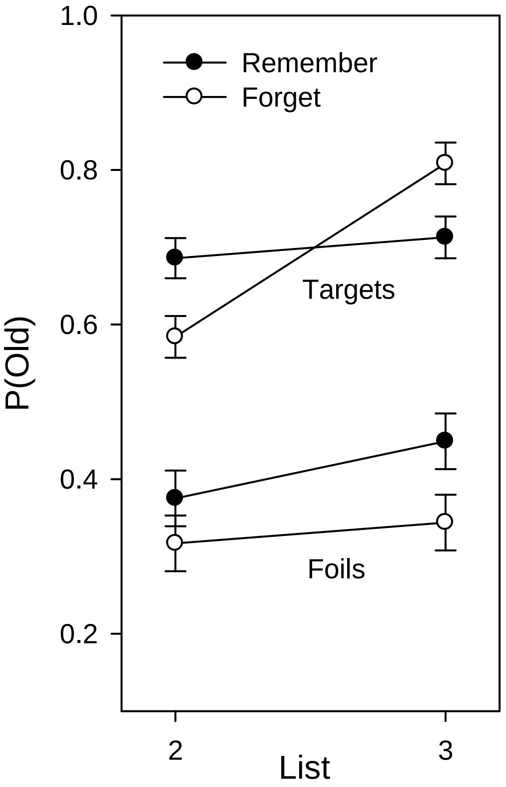

Note. In this experiment, participants were to only endorse items that were studied on a specific list. Thus, some items should have been rejected even though they were studied because they were studied on the to-be-excluded list and not studied on the to-be-endorsed list. This graph shows hit rates for list 2 and list 3 in the remember and the forget conditions (targets), along with false-alarm rates for items that were studied on the to-be-excluded list (foils). For example, if the participant was instructed to positively endorse only items on list 2, any list 3 items that were positively endorsed counted as list 3 foils.

Exclusion Results

Hits. For the Exclusion condition, participants were asked to recognize either from L_2 or L_3, so the data was analyzed as a two-way ANOVA with both List and Instruction as between-subject factors. The hit rates were greater for L_3 than L_2, $F(1,144) = 22.06$, $MSE = .026$, $p < .001$. While there was no significant main effect of Instruction, there was a significant List x Instruction interaction, *F* = 13.54, *p* < .001. As shown in *Figure* 4, I found a crossover interaction, just as in Experiment 1, with hit rates in the remember condition greater than in the forget condition for L_2, and the opposite for L_3. Planned comparisons confirm the simple effects. Thus in terms of hit rates, I found the costs and benefits of directed forgetting, and a recency effect. As displayed in *Figure 5*, there were no significant effects of serial position, *F*(15,144) = 1.05, *MSE* = .225, *p* = .399.

Figure 5. Exclusion serial position data (Experiment 3).

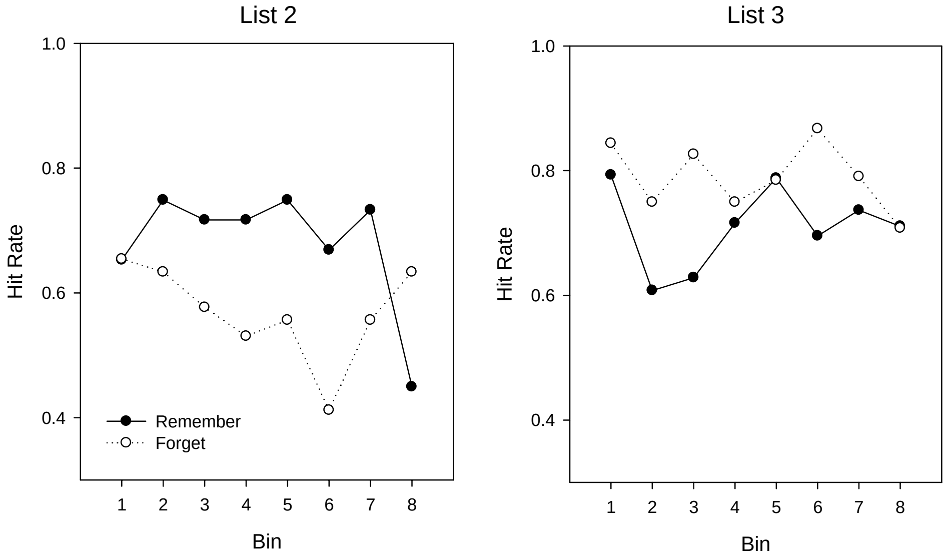

Note. For the sake of clarity, the 16 item list was compiled into 8 bins spanning two serial positions. For instance, bin *n* contains the data from serial positions *2n-1 and 2n.*

False Alarms. First let us consider the false-alarm rates for those test items that were not studied. These are shown in *Table 1*. Analysis of the false alarm rates for new words shows no significant main or interaction effects. This suggests that all things being equal, recognition performance (whether an item was studied or not) is captured solely by the differences that were observed in hit rates. The cross-over interaction that was observed in hit rates, therefore, indicates that there are costs and benefit associated with the instruction to forget for recognition memory just as is the case for free recall.

The exclusion instructions were to respond negatively to items that were studied but were not studied on the target list. The rate at which participants failed to do so can also be considered a false-alarm rate. There was significant main effect of instruction on the false-alarm rates, $F(1,144) = 5.14$, $MSE = .049$, $p = .03$. As shown in *Figure* 4, false-alarm rates are uniformly lower for participants in the forget condition than in the remember condition. This is consistent with the context model insofar as an acceleration of the change of context between the two lists as the result of the instruction to forget should make it easier to reject items from the wrong, non-target list. Additionally, while the effect of List is not reliable, participants responded "yes" to L_3 items when being tested on L_2 more often than they responded "yes" to L_2 items while being tested on L_3. This trend is consistent with the assumption that the context used to probe memory is more similar to L_3 than to L_2. Thus, more L_3 items are mistakenly associated with the L_2 context when memory is probed and the target list is L_2 and vice versa when L_3 is the target list.

Inclusion Results

The inclusion version of the recognition experiment asked participants to recognize words that had been studied on both lists 2 and 3 – thus it was analyzed in a mixed-model ANOVA with List as a within-participants variable and Instruction as a between-subjects variable. There was a significant main effect of List, $F(1,58) = 35.43$, $MSE = .018$, $p < .001$; hit rates were higher for L_3 than for L_2.

Figure 6. Inclusion recognition performance (Experiment 3).

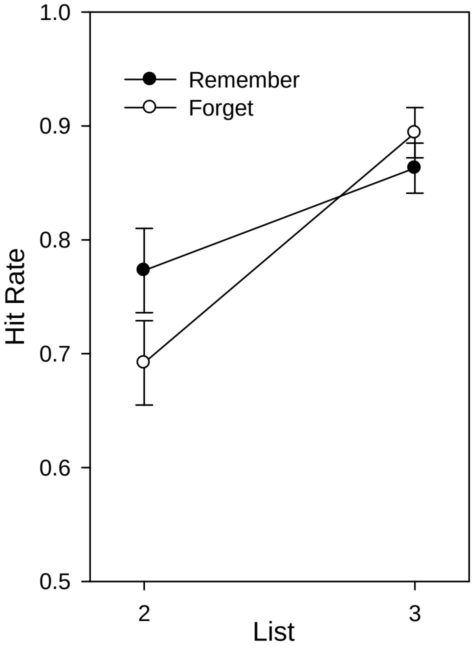

Just as in recall, there was no main effect of Instruction, but there was a significant List x Instruction interaction, *F*(1,58) = 5.271, $p = .025$. As shown in *Figure* 6, hit rates were higher for the remember condition on L_2 but there was a minimal difference in recognition on L_3. Planned comparisons revealed that performance was significantly higher for L_2 in the remember condition, but the difference between conditions was not significant for L_3 ($p = .47$). There were no differences in probability of responding "yes" to new words between remember and forget conditions, as shown in *Table 1*. As shown in *Figure* 7, there was not a significant effect of serial position in the inclusion experiment, *F*(15,58) = 1.43, *MSE* = .147, $p = .123$.

Figure 7. Inclusion serial position data (Experiment 3).

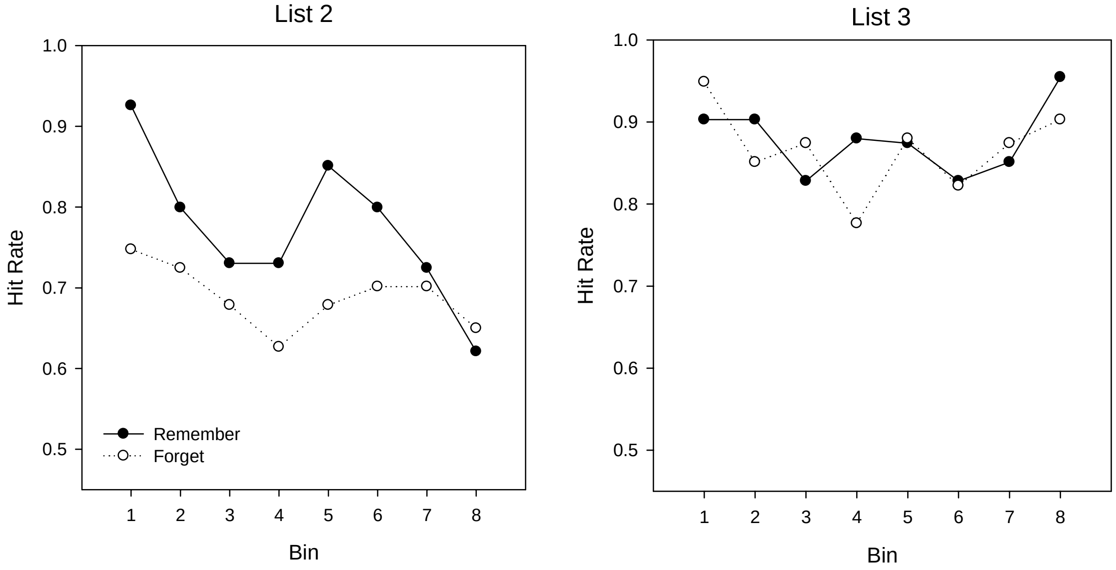

Note. For the sake of clarity, the 16 item list was compiled into 8 bins spanning two serial positions. For instance, bin *n* contains the data from serial positions *2n-1 and 2n.*

Discussion

The results of Experiment 2 are clearly inconsistent with the conventional wisdom of the field. I observed clear costs and benefits for recognition memory as a result of the instruction to forget. Moreover, Experiment 3 shows that there are costs and benefits associated with the instruction to forget even for an inclusion memory task. The costs appear to be larger than the benefits. However, the inclusion task is easier than the exclusion task, and thus hit rates are relatively high, especially for L_3. Thus, I was concerned that the relatively small benefits might be due to some participants performing at ceiling. This concern motivated the next two experiments.

Experiments 4 and 5

Because of the potential of ceiling effect masking the benefits of directed forgetting in Experiment 3, Experiment 4 is a replication of Experiment 3 with a reduced study time. In order to replicate my findings and to allow for a further test of the model, Experiment 5 is a replication of Experiment 2 (Exclusion) also with a reduced study time.

Methods

Participants, Materials, and Procedure. In Experiment 4 an inclusion procedure was used. Participants were 86 (43 in each condition) undergraduate psychology students at the University of South Florida who participated in exchange for course credit. The procedure was identical to that of Experiment 3, but with a 4 second study time.

In Experiment 5, an exclusion procedure was used. Participants were 172 (43 in each condition) undergraduate psychology students at the University of South Florida who participated in exchange for course credit. The procedure was identical to that of Experiment 2, but with a 4 second study time.

Inclusion Results

As in Experiment 3, the data was analyzed in a mixed-model ANOVA with List as a within-participants variable and Instruction as a between-subjects variable. There was a significant main effect of List, $F(1,82) = 45.49$, $MSE = .019$, $p < .001$; hit rates was higher for L_3 than for L_2. There was no main effect of Instruction, but there was a significant List x Instruction interaction, $F(1,82) = 7.21$, $MSE = .019$, $p = .009$. As shown in *Figure 8*, hit rates were better for participants in the remember condition on L_2 and the forget condition on L_3. Planned comparisons revealed that all differences shown here were significant. As predicted, I saw recency of L_3 and both the costs and benefits of directed forgetting. As in Experiment 3, I saw no difference in false alarm rates between the remember and forget conditions (see *Table 1)*. Because there were no significant serial position effects in Experiments 2 and 3, serial position effects were not examined for Experiments 4 and 5.

Figure 8. Inclusion 4s recognition performance (Experiment 4).

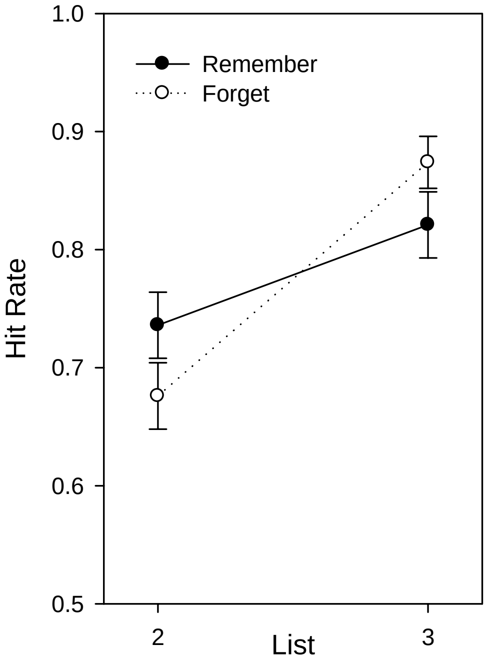

Exclusion Results

Hits. As in Experiment 2, the data was analyzed as a two-way ANOVA with both List and Instruction as between-subject factors. The hit rates were greater for L_3 than L_2, $F(1,168) = 15.92$, $MSE = .028$, $p < .001$. While there was no significant main effect of Instruction, there is a significant List x Instruction interaction, $F = 6.71$, $p < .001$. As shown in *Figure 9*, I found a crossover interaction, with hit rates in the remember condition greater than in the forget condition for L_2, and the opposite for L_3. Planned comparisons confirm the simple effects. In terms of hit rates, I replicated the pattern of Experiment 2.

Figure 9. Exclusion 4s recognition performance (Experiment 5).

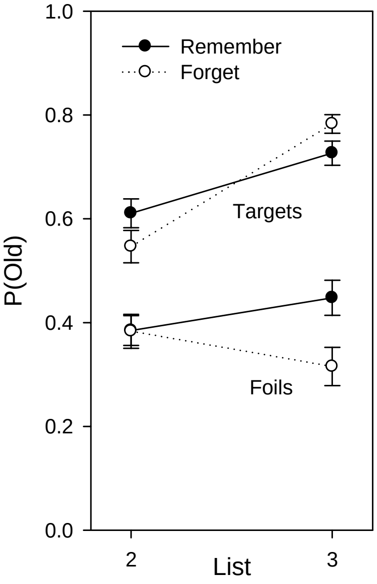

Note. In this experiment, participants were to only endorse items that were studied on a specific list. Thus, some items should have been rejected even though they were studied because they were studied on the to-be-excluded list and not studied on the to-be-endorsed list. This graph shows hit rates for list 2 and list 3 in the remember and the forget conditions (targets), along with false-alarm rates for items that were studied on the to-be-excluded list (foils).

False Alarms . False alarm patterns for those test items that were not studied (unstudied foils) are similar to those in Experiment 2; however the differences between lists are now significant, $F(1,168) = 21.12$, $MSE = .029$, $p < .001$. These are displayed in *Table 1*. In the exclusion experiments, the rate at which participants responded positively to items that were studied but not on the target list (foils) is also a false-alarm rate. In the exclusion experiment with a shortened study time, some differences emerge compared with the longer study time. There is no longer a main effect of instruction, as there was in the long study time version, but there is now a significant List x Instruction interaction, $F(1,168) = 7.55$, $MSE = .047$, $p = .007$. As shown in *Figure* 9, the advantage of lower false-alarm rates for the forget condition is only apparent on L_3; there is no difference in false-alarm rates on L_2.

Discussion

In both inclusion and exclusion with reduced study time, both costs and benefits were found. The shortened study time was successful in pulling participants away from ceiling and allowing significant benefits of directed forgetting to emerge. Contrary to previous findings in recognition findings, we found an effect of directed forgetting in both free recall and recognition. Findings from the recognition experiments are consistent with our context change model of directed forgetting.

Experiment 6 – Delayed Free Recall

If the costs and benefits involved in directed forgetting are due to a context change, then increasing the delay between study and test should change the context, such that it is less similar to the context of L_2 or L_3. Thus, I expect to see a decrease in the recency effect that I observed in Experiment 1. Moreover, those features that discriminate L_2 and L_3 should be more difficult to reinstate after a relatively long delay. If this is the case, then both the costs and the benefits of directed forgetting should be eliminated according to the present model. The results obtained from this experiment will help to understand the nature of directed forgetting. One question the data will address is whether the effect of directed forgetting is to instigate a permanent change in the state of a memory trace, which is implied by many different hypotheses.

Method

Participants, Materials, and Procedure. Participants were 176 (44 in each condition) undergraduate psychology students at the University of South Florida who participated in exchange for course credit. The procedure was identical to that of Experiment 1 (free recall) experiment, except with an increased lag between study and test. After study of the third list (and completion of the third distractor task), participants engaged in a 5 minute delay task, designed to prevent rehearsal, after which the test appeared exactly as it did in Experiment 1.

Delay task. After completion of the third distractor task, participants were told that they would next see a short video and then they would be tested on the video. In order to encourage attending to the video and prevent rehearsal of words from the lists, participants were told that they must pass the quiz in order to continue to the next phase of the experiment (although this was not enforced). They then watched a four and a half minute informative video about how contact lenses are made, followed by a five question quiz that took approximately 30 seconds. After the quiz, they went on to the free recall test, which took place exactly as in Experiment 1.

Results

Correct Recall. The results of a two-way ANOVA show that for correct recall, there was a main effect of List, $F(1,172) = 10.05$, $MSE = .024$, $p < .001$, and a main effect of Instruction, $F(1,172) = 6.27$, $p = 013$; probability of recall was greater for L_3 than for L_2, and greater for the forget condition than the remember condition. Both main effects are qualified by a significant List x Instruction interaction, *F*(1,172) = 6.58, *p* = .011. As shown in the left-panel of *Figure 10,* there was no difference in recall for L_2, but recall was higher in the forget condition than the remember condition on L_3.

Figure 10. Delayed free recall performance (Experiment 6).

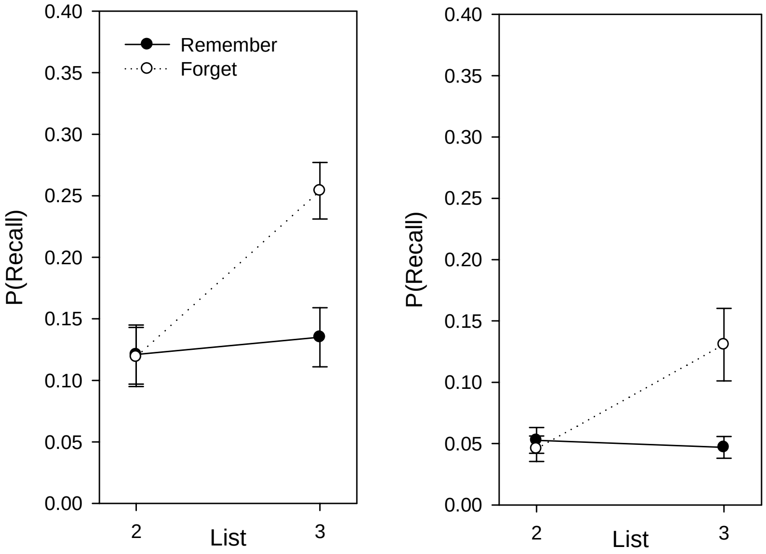

Note. The left panel shows probability of recall for correct items. The right panel shows probability of intrusions that came from either list 2 or list 3. When recalling from list 2, any list 3 item that was output is referred to as an intrusion and vice versa.

Serial Position. In delayed free recall, significant interaction effects were found for serial position, *F*(15,172) = 3.290, *MSE* = .057, *p* < .001. As shown in *Figure 11,* serial position curves did not differ for L_2; however the primacy effect is significantly larger for L_3 in the forget condition than in the remember condition. While there was a significant advantage for the beginning of the list in first item output, there was not a significant interaction between first item output and instruction, *F*(15,172) = 1.491, *MSE* = .125, *p* = .10.

Figure 11. Serial position in delayed free recall (Experiment 6).

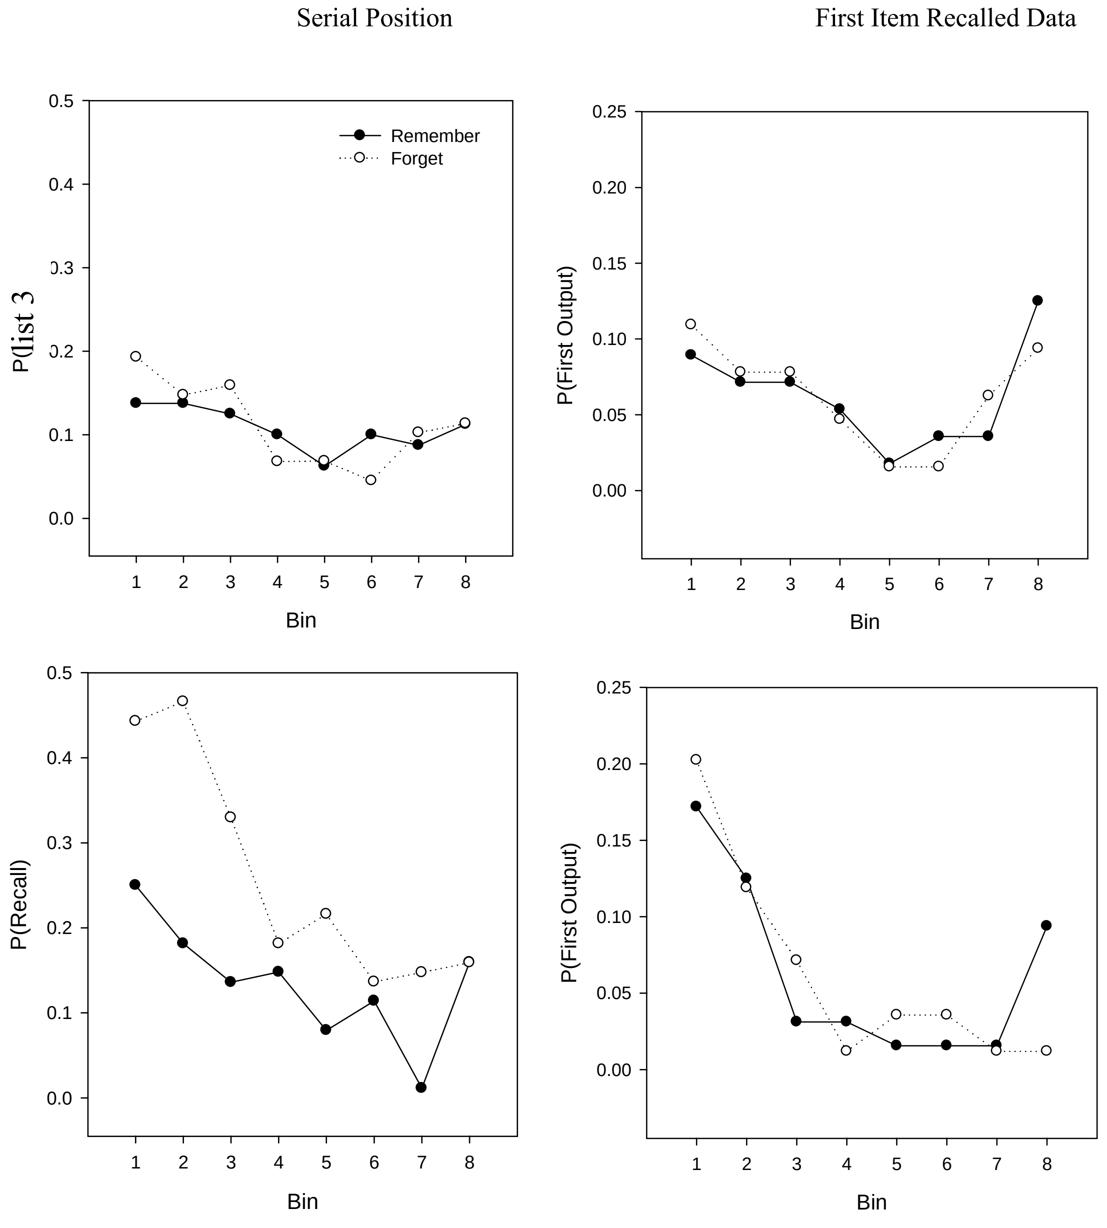

Note. For the sake of clarity, the 16 item list was compiled into bins. For serial position, each bin spanned two serial positions. For instance, bin *n* contains the data from serial positions *2n-1* and *2n.* For first item output, bin 1 represents the first item on the list (since this is where differences are seen) and all other serial positions are grouped by three.

Unlike in immediate free recall, there were no significant differences in CRP between lists, $F(1,90) = 1.073$, $MSE = .077$, $p = .301$. As shown in *Figure 12*, on both the L_2 and L_3, probability of a forward lag of 1 is greater for the remember than the forget conditions.

Figure 12. Conditional response probabilities for delayed free recall (Experiment 6).

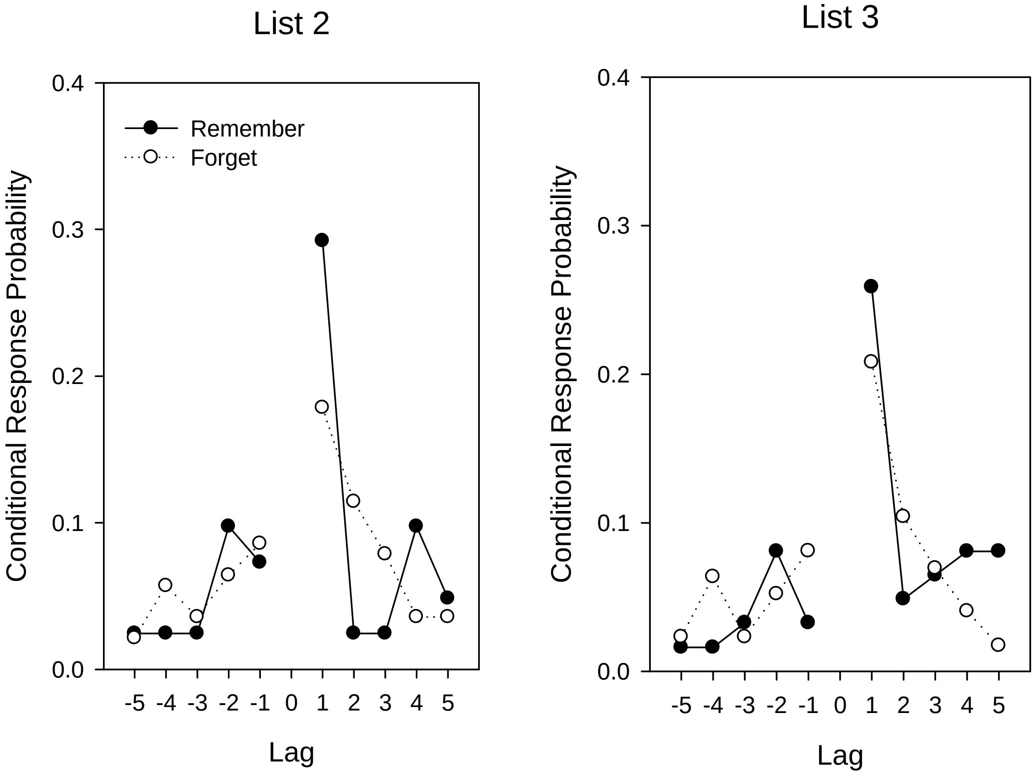

Note. Conditional response probabilities refer to the probability of recalling an item a given distance away from the previous item in serial position. Successive recall of nearby items is more likely than recall of distant items, and forward movement is more likely than backward.

Intrusions. Compared with free recall without the five-minute delay, recall after the five minute delay also leads to a different pattern of intrusions. As shown in the right panel of *Figure 10*, intrusions from any list are higher in this experiment compared with the no-delay free recall. A higher level of intrusions suggests that the context that differentiates the lists is not being used to distinguish the lists after the delay, leading to confusion as to which list the words came from.

For intrusions from L_2 when recalling L_3 and vice versa, there was a significant main effect of List, $F(1,172) = 5.38$, $MSE = .013$, $p = .022$, a main effect of Instruction, $F(1,172) = 5.10$, $p = .025$, and a List x Instruction interaction, $F(1,172) = 7.04$, $p < .001$. The pattern of intrusions exactly replicates the correct recall pattern. There was no difference between remember and forget conditions for L_2 intrusions when recalling L_3, but there were significantly more intrusions from L_3 when recalling L_2 in the forget condition than in the remember condition. Further, there was no difference in recall between L_2 and L_3 for the remember condition.

For intrusions that came from L_1, there is a significant main effect of List, $F(1,172) = 9.78$, $MSE = .006$, $p = .002$. As expected, participants were more likely to have intrusions from L_1 while they were recalling L_2 than when they were recalling L_3 for both remember and forget conditions. Additionally, there was a significant List x Instruction interaction, $F(1,172) = 4.23$, $p = .04$. The difference in L_1 intrusions between L_2 and L_3 was greater in the remember condition than in the forget condition. L_1 intrusions are displayed in *Table 1*.

Discussion

The delayed free recall experiment supported the context-change hypothesis. First, it eliminated the recency effect of L_3 seen in the immediate free recall experiment. Because of the delay between study and test, the context during test is no longer a very good match to the L_3 context, thus its recency advantage disappeared. In addition, the costs were eliminated, as expected, but not the benefits of directed forgetting. Costs are diminished because a large change in context eliminates the advantage of L_2 in the remember condition, and hence the text context is not very similar to either list. Additionally, when comparing these findings to those in from Experiment 1, one can see that recall is better on L_2 in the forget condition after the delay. This suggests that a reduction in L_3 recall (also apparent after the delay) allows for less interference from L_3 items, thus more sampling of L_2.

While it might be expected that the delay would also eliminate the benefits, the results can be explained on the assumptions that the delay does not negatively affect the ability to reinstate the most recent L_3 context as much as it does the more distant L_2 context. It does appear in the lower right panel of *Figure 11*, however, that subjects have a more difficult time initially retrieving the first item on L_3 after the delay. After the delay, the first-item recall function is much shallower when compared to the same function from Experiment 1. Thus, subjects were more likely to initiate retrieval in the later serial positions. This was even more so the case for L_2 where subjects often initiated retrieval by outputting from the end of the list.

According to Sahakyan and Delaney's (2003) differential encoding hypothesis, we would expect the elimination of the costs but not the benefits. However, that hypothesis is not supported when the full set of data are considered. Sahakyan and Delaney hypothesized that while the costs are due to a context change, the benefits are due to a change to a better encoding strategy between lists that occurs only in the forget condition. If so, we would expect to observe benefits for all serial positions on L_3, but we only observe strong benefits for items on the first half of L_3 in much the same we observed benefits after only short delay in Experiment 1.

In addition to correct recall, the intrusion rates seen in the delayed free recall experiment are also consistent with both a context change and a differential encoding hypothesis. First, the intrusion rates were higher in delayed free recall compared to immediate free recall, suggesting that the context that participants use to differentiate the lists is not as readily available after a delay; thus when an item from the wrong list is sampled, participants are less able to judge that this came from the wrong list based on the different contexts of the two lists. The intrusion rates for items that came from L_2 or L_3 exactly mirrored the correct recall patterns of data – with significantly more intrusions from L_3 when trying to recall from L_2 in the forget condition compared to all other conditions. While in free recall, we expected to see an overall lower intrusion rate for the forget condition, this is not the case in the delay experiment. After the five minute delay, the context has changed such that participants are no longer as good at using context to differentiate lists and judge whether an item came from the wrong list. The higher output of L_3 items in the forget condition suggests that these items are more accessible due to stronger encoding; since participants are not using the context that differentiates the lists, these highly activated items appear as both correctly recalled items and as intrusions.

Chapter Three: A Formal Model of Directed Forgetting

I will now discuss a formal model of the context change that occurs with directed forgetting. The model must be able to account for findings in both free recall and recognition. It must be able to handle not only the costs and benefits of directed forgetting and the recency of L_3, but also the serial position curves, first-item output, and CRP functions. The first-item output findings are particularly important because most of the directed forgetting effect in free recall appears to be driven by differences in the first-item output patterns. Finally, the model must also be able to account for recognition with shortened study time and for delayed free recall.

A REM Model

I have developed a model in the REM framework that accounts for all of the directed forgetting data mentioned above in terms of the context change explanation by extending the REM free recall and context encoding model (Malmberg and Shiffrin, 2005) and the REM recognition models (Malmberg, Holden, & Shiffrin, 2004; Xu & Malmberg, 2007; Shiffrin and Steyvers, 1998). The results are a single set of well-specified assumptions that account for the effect of directed forgetting on free recall, inclusion recognition, and exclusion for both list method and the item method procedures.

Representation

According to REM, general knowledge of items is stored in lexical/semantic memory traces and information about past events is stored in episodic memory traces. Lexical/semantic traces are acquired over a lifetime. They contain information about how words are spelled and pronounced and what they mean. In addition, they contain information about the contexts or situations in which they have been encountered. As such, they are accurate, complete, and generalizable to the contexts in which they usually occur. Two concatenated vectors of features represent these traces. One vector represents the item and the other represents the contexts in which it has been encountered. These vectors are generated according to a geometric distribution with the base rate parameter, *g*:

$$P[V = j\} = (1-g)^{j-1} g, j = 1,...,\infty . \qquad (1)$$

When a word is studied, the w_i item features of its lexical semantic trace are copied to form a new episodic trace that represents this occurrence. In addition, w_c features of the current context are stored. Episodic encoding is an incomplete and error-prone process. During the storage process, a feature may be copied correctly, it may be copied incorrectly, or it may fail to get copied at all. The probability of storing a feature given a certain unit of time (*t*) is represented by the u_x^* parameter. Given that a feature is stored, it is stored correctly with a probability *c*. An item will be stored incorrectly with a probability *1-c*, in which case a feature will be randomly chosen according to the geometric distribution. The absence of a stored feature is represented by the value zero.

When items are studied, context is stored in episodic traces in the same way. For the sake of simplicity, I will assume that context features change between lists with a probability of β, but not within lists. Thus, all items within a list will share the same context information. When a context feature value is changed it is randomly sampled from the geometric distribution. I further assume that context features change after the final study list. The context changes in the same manner between the final list and test. The features representing the item itself, however, will be different for each item.

Buffer Operations

As a descendent of the modal model (Atkinson and Shiffrin, 1968), the interaction of control processes and structural aspects of memory are used to model serial position data. Control processes operate on items located in a limited capacity rehearsal buffer during encoding. For present purposes, I chose a buffer capacity of two items, although the larger capacities would also work. Upon the presentation of the first item on a list, its lexical/semantic item features enter the buffer and two things happen. First, t_o attempts to copy the items features in an episodic trace are made. The probability of storing an item feature is u_i^*. Second, t_1 attempts to copy the current context features in an episodic trace are made. The probability of storing a context feature is u_c^*. I assume that attention is focused on the item itself rather than context, so item information will be better encoded than context information, represented by a greater u^* value for item information ($u_i^* > u_c^*$).

Upon the presentation of the second item on the list, its lexical/semantic item features take the remaining slot in the buffer, and now three things happen. The item and context features are stored in the same way as before. However, in addition, some of the items features of the first item are stored in another concatenated vector (cf. Kimball, Smith & Kahana, 2007). This represents the assumption that an episodic association between the two items in the buffer is stored (this loosely corresponds to strengthening an inter-item association in SAM). The probability of storing the associative items features is u_a^*. I assume that attention is primarily focused on the present item; information about the older item in the buffer will be encoded worse than the current-item information, represented by a greater u^* value for item information ($u_i^* > u_a^*$). That is, I assume that participants will tend to focus their attention on the most novel items in the rehearsal buffer. As new items are added to the buffer, this process repeats, with the oldest item being dropped with a probability δ.

Retrieval

The first step of the retrieval process is similar across all test conditions (recall, recognition-inclusion, recognition-exclusion). An activated subset is created, which consists of only the items with the strongest association to the current context. From a participant's perspective, only items encountered recently, that is during the experiment, are relevant to the task. In order to access only these relevant items, the retrieval task uses only items in this subset.

In order to create the active subset, the current context cue is matched against the episodic images stored in memory. The matching process involves calculating a likelihood ratio for each trace, which takes into account both features that match and features that do not match. Matching features increase and mismatching features decrease the likelihood ration; cases where no features are stored do not contribute to the likelihood ratio either way. Likelihood ratios are calculated according to the following equation:

$$\lambda_j = (1-c)^{n_{ij}} \prod_{i=1}^{\infty} \left[\frac{c + (1-c)g(1-g)^{i=1}}{g(1-g)^{i-1}} \right]^{n_{ijm}}, \qquad (2)$$

where g is the environmental base rate for the occurrence of features, i is a feature value ranging from 1 to infinity, n_{ij} is the number of mismatching context features for an item (regardless of their value), and n_{ijm} is the number of times feature i matched the retrieval cue with value j. The "activated" subset of memory will consist of a certain percentage of all traces in memory, represented by the ρ parameter. The items that get into the subset are those with the highest likelihood ratios.

Free Recall

The free recall task begins with the creation of the cue with which to probe memory. The initial cue consists of only context features. I further assume that the context cue is a combination of the current test context and reinstated list context, in order to allow one to access the intended list. The proportion of reinstated list context features is represented by the γ parameter. Free recall operates in REM as a memory search process, with cycles of sampling and recovery (Malmberg & Shiffrin, 2005). For simplicity, I assume that if an item is sampled, it is recovered with a probability of 1.0. The cue is matched against all traces in the subset in an attempt to sample an item from the given list. Likelihood ratios for all images are calculated according to Equation 2. The probability of sampling image, I_i, given the context retrieval cue, Q, is as follows

$$P(I_i \mid Q) = \frac{\lambda_i}{\sum \lambda_k} \tag{3}$$

If an item is sampled and it comes from the correct list, it is output with a probability of 1.0. If, however, an item is sampled and it comes from an incorrect list, I assume that the participant undergoes a monitoring process. The probability of outputting an incorrect item is a function of the overlap in context between lists (represented by the η parameter). If the context change that occurs between lists is large, then there is little overlap between lists and participants will be better able to judge whether a sampled item came from the correct list or not. If the context change between lists is small, there will be much overlap between the lists, and it will be much harder for participants to judge whether a sampled item came from the intended list.

If an item is output, the next cue used to probe memory will consist of both context and item information. Again, the context portion of the cue consists of both current context features and context features associated with the given list. The item portion of the cue consists of the item vector from the last item recalled. Thus it is most likely that co-rehearsed items, which share the current item's information, will be sampled next. If no item is output, then the original context cue is used for the next probe of memory. The sample and recovery process repeats κ times.

Recognition – Inclusion

In the inclusion task, a participant's task is to positively endorse any studied items. For this reason, a simple global matching process is used. In REM, a decision about whether an item is judged as "old" is made based on the likelihood ratios calculated for all items in the comparison set. In this case, the cue (a test item) is compared to all items in the activated subset. The "odds" are calculated according to the following equation:

$$\Phi = \frac{1}{n}\sum_{j=1}^{n}\lambda_j \ , \tag{4}$$

and if the odds exceed a specified criterion, the item is judged as old, otherwise it is judged as new. In the absence of any instructions that would lead to a bias to respond differently, an old decision is made if the odds exceed a criterion of 1.0, thus this is the decision criterion that I use here.

Recognition – Exclusion

In the exclusion task, a participant's task is to positively endorse only items that came from a given study list. Foils consist of both new items and items from the other list, so participants must be able to distinguish between studied items from the correct list and studied items from the other list. A global matching process, as that in the inclusion task, is first used followed by a monitoring task, as in the free recall task. After an item is identified as a studied item based on the global match (according to Equation 4), a participant again makes an output decision that is dependent on the overlap in context between the two lists. Large overlap in context means that it is harder to distinguish between the two lists and the false alarm rate will be increased.

Effects of the Forget Instruction

The forget instruction will have multiple effects in these tasks. First, it will increase the rate of context change between lists, so that the two lists share less context features (less overlap). Second, it will alter the encoding of the first item on the third list in the forget condition. The first item on any list is encoded more strongly than other items because the rehearsal buffer is not full during encoding of this item, and it has more opportunities to be linked to the context of the list. I assume that the opportunity for linking this item to context is greater for the third list in the forget condition, because this list is encoded after the forget instruction, when no items from the previous list are being rehearsed. Finally, the forget instruction will decrease the probability of reinstating features for use in the context cue used in the recall process. Because the forget instruction increases the context change that occurs between lists, it should be harder to reinstate context features after this instruction. As mentioned previously, the probability of an intrusion error will be dependent on the context change that occurs between lists, thus the probability of outputting an item from the wrong list will be lower in the forget condition.

I attempted to find a reasonable set of parameters to account for my more than 300 data points. I did not attempt to find a "best-fitting" set of parameters, and I am focused more on accounting for the overall patterns of the observed data, and less concerned with formal model comparisons. That said, I did vary a number of parameters in order to determine if the quantitative predictions were in the "right ballpark". Descriptions of each parameter are listed in Table 2. Many parameter values are common to all experimental conditions. In addition, there are 11 free parameters, but the majority of these are scaling parameters that do not differ between conditions; only four parameters differ between remember and forget conditions and these are carry most of water for the model. The parameters that will differ from those listed in the table after the forget instruction are as follows: $t_1 = 12$; $\beta = .8$; $\rho_2 = .15$; $\eta = .5$.

Data and model predictions for Experiments 1 through 3 are presented in *Figure*s 13-20. Overall there is a strong correspondence between the model and data. The model captures the costs and benefits of directed forgetting for free recall, exclusion recognition, and inclusion recognition. It also accounts for the serial position, first-item output, and CRP data for free recall. While a more complicated model could probably do a better job than the current model, there does not seem to be much to be purchased by the additional complexity.

Table 2.

Parameter Values and Descriptions

Parameter	Value	Description
g	.4	Environmental base rate (standard value)
w_i	8	Number of item features
w_c	8	Number of context features
c	.8	Probability of correctly storing a feature
u^*_i	.5	Probability of storing an item feature
u^*_c	.2	Probability of storing a context feature
u^*_{cr}	.1	Probability of copying a co-rehearsed item's feature
t_1	6*	Number of storage attempts for first item on a list
t_0	2	Number of storage attempts for all other items on a list
κ	20	Number of sampling attempts
β	.2*	Probability of change for context features between lists
δ	.75	Probability of dropping the oldest item in the buffer
ρ_2	.2*	Probability of reinstating context features on list 2
ρ_3	.8	Probability of reinstating context features for list 3
σ	.8	Size of activated subset of items
η	.6*	Probability of outputting an intrusion

Note. Parameter values with asterisks are those that differ in the forget condition. For the forget condition, $t_1 = 12$; $\beta = .8$; $\rho_2 = .15$; $\eta = .5$.

Figure 13. Model predictions for correct recall and intrusions in free recall.

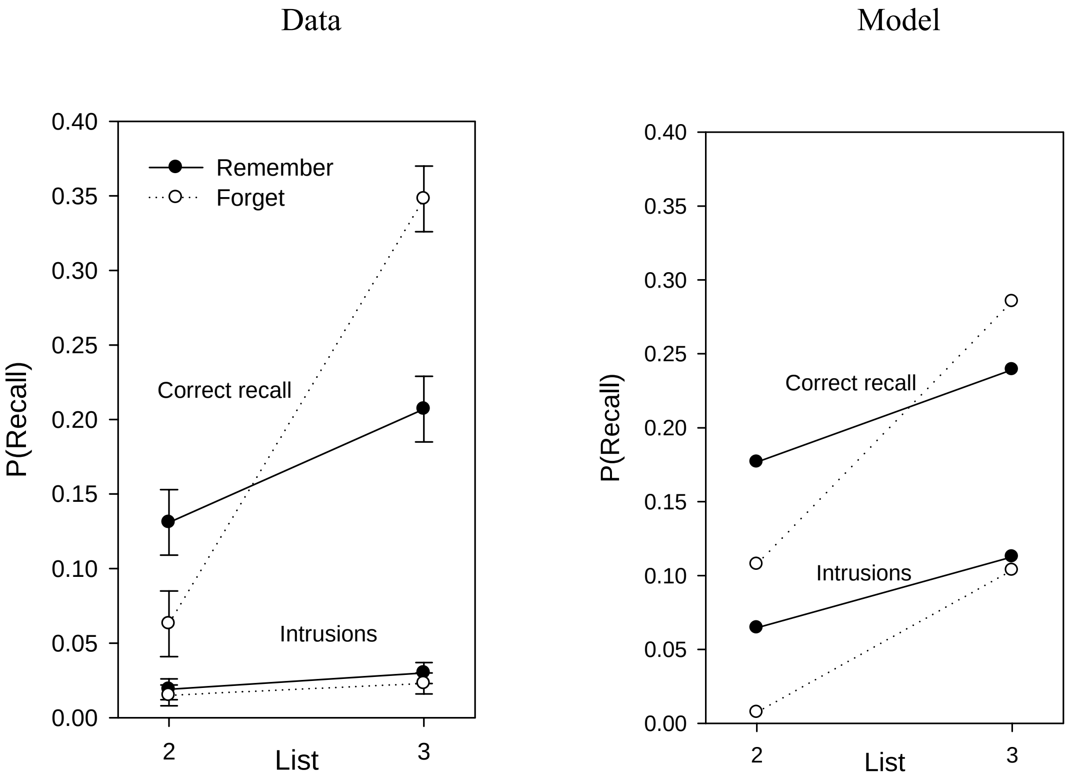

Figure 14. Model predictions for serial position data in free recall.

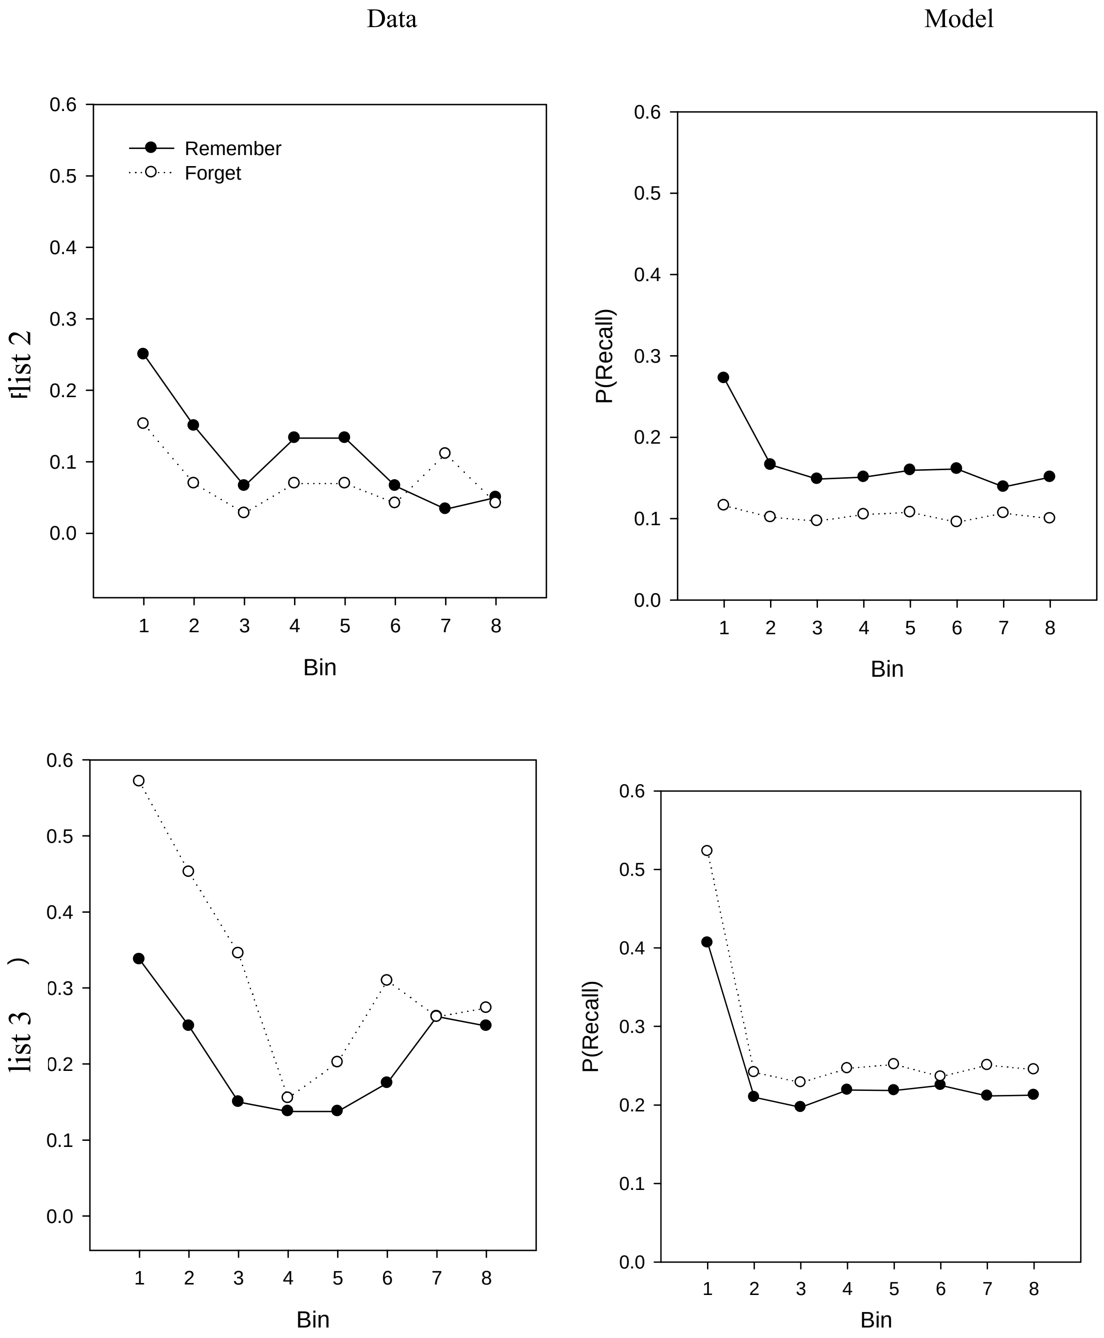

Note. For the sake of clarity, the 16 item list was compiled into 8 bins spanning two serial positions. For instance, bin *n* contains the data from serial positions *2n-1 and 2n.*

Figure 15. Model predictions for first item output position data in free recall.

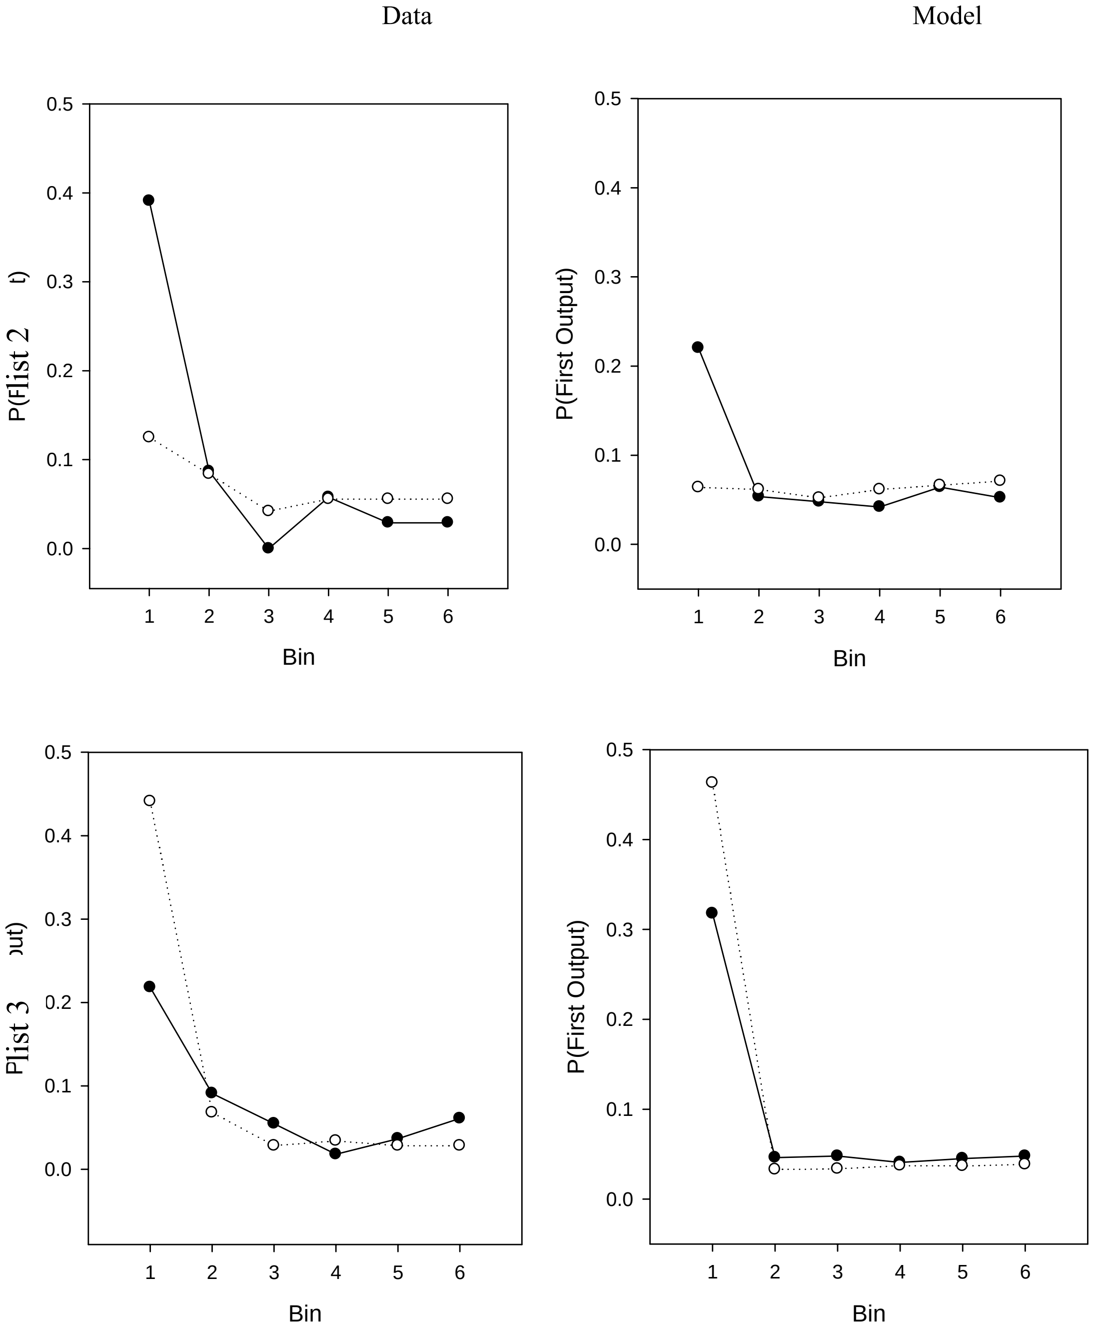

Note. For the sake of clarity, the 16 item list was compiled into bins. For first item output, *bin 1* represents the first item on the list (since this is where differences are seen) and all other serial positions are grouped by three.

Figure 16. Model predictions for conditional response probabilities from free recall.

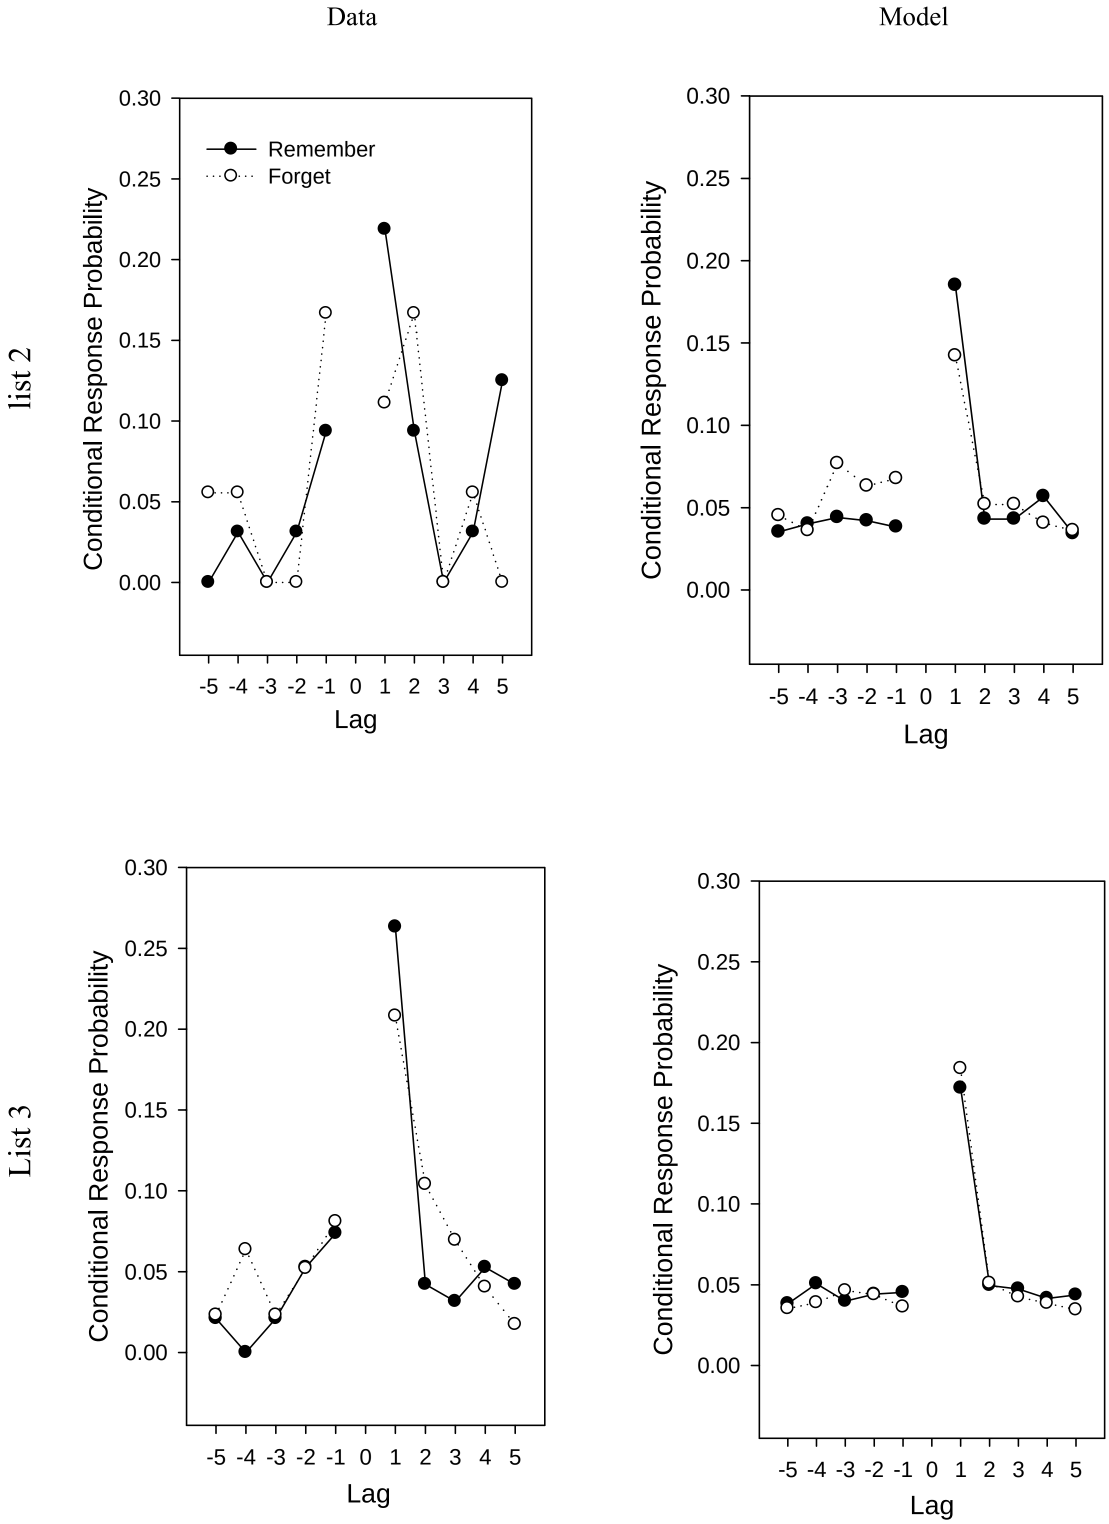

Figure 17. Model predictions for inclusion recognition.

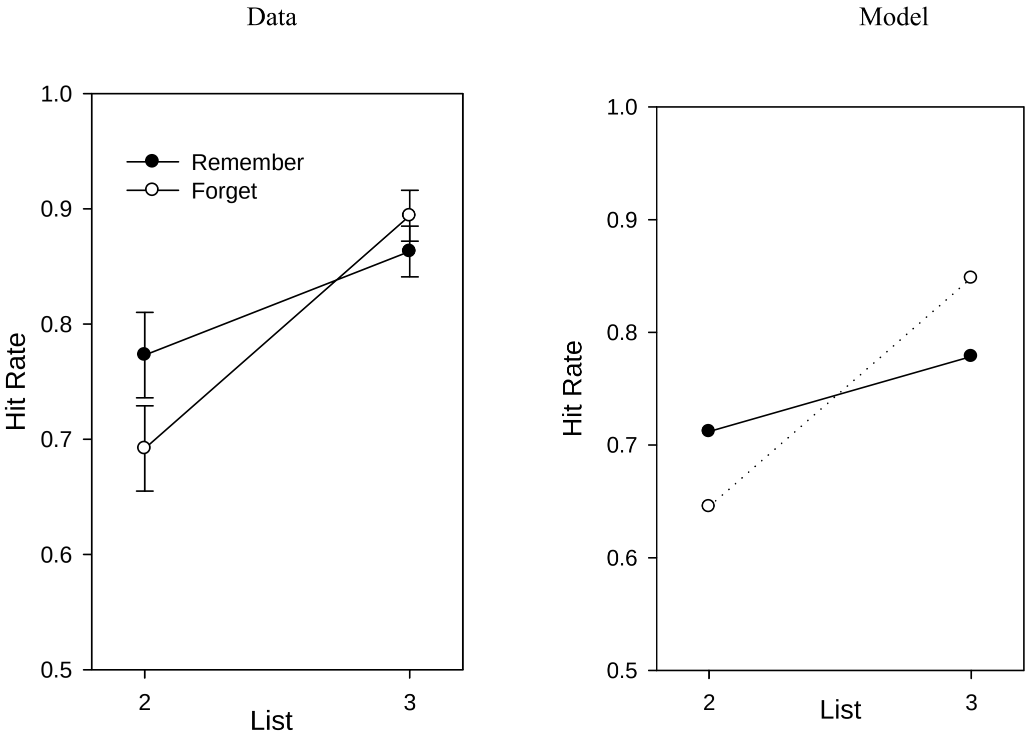

Figure 18. Model predictions for serial position data in inclusion recognition.

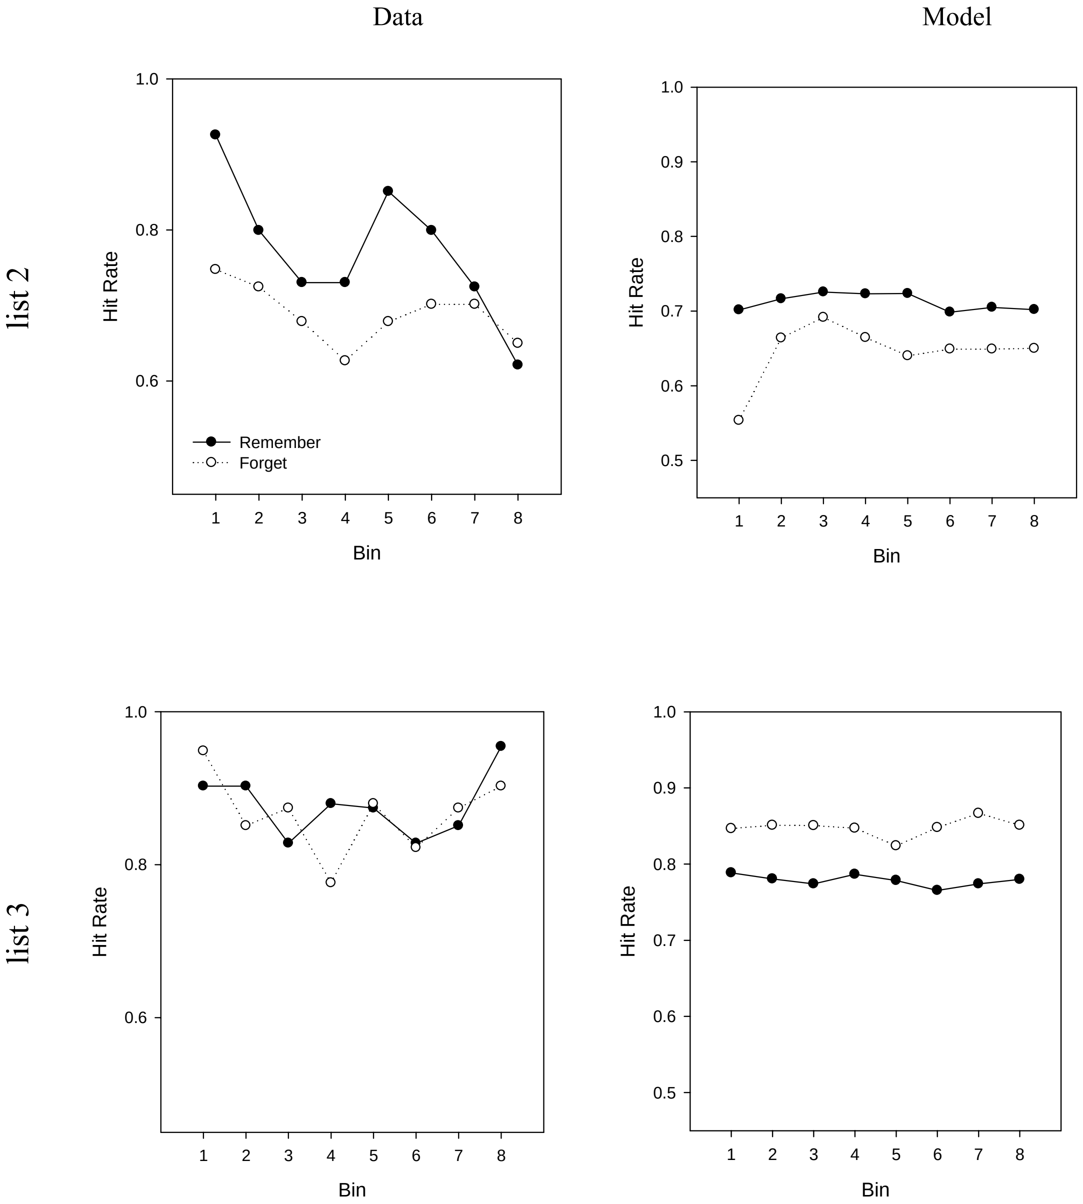

Note. For the sake of clarity, the 16 item list was compiled into 8 bins spanning two serial positions. For instance, bin *n* contains the data from serial positions *2n-1 and 2n.*

Figure 19. Model predictions for exclusion recognition.

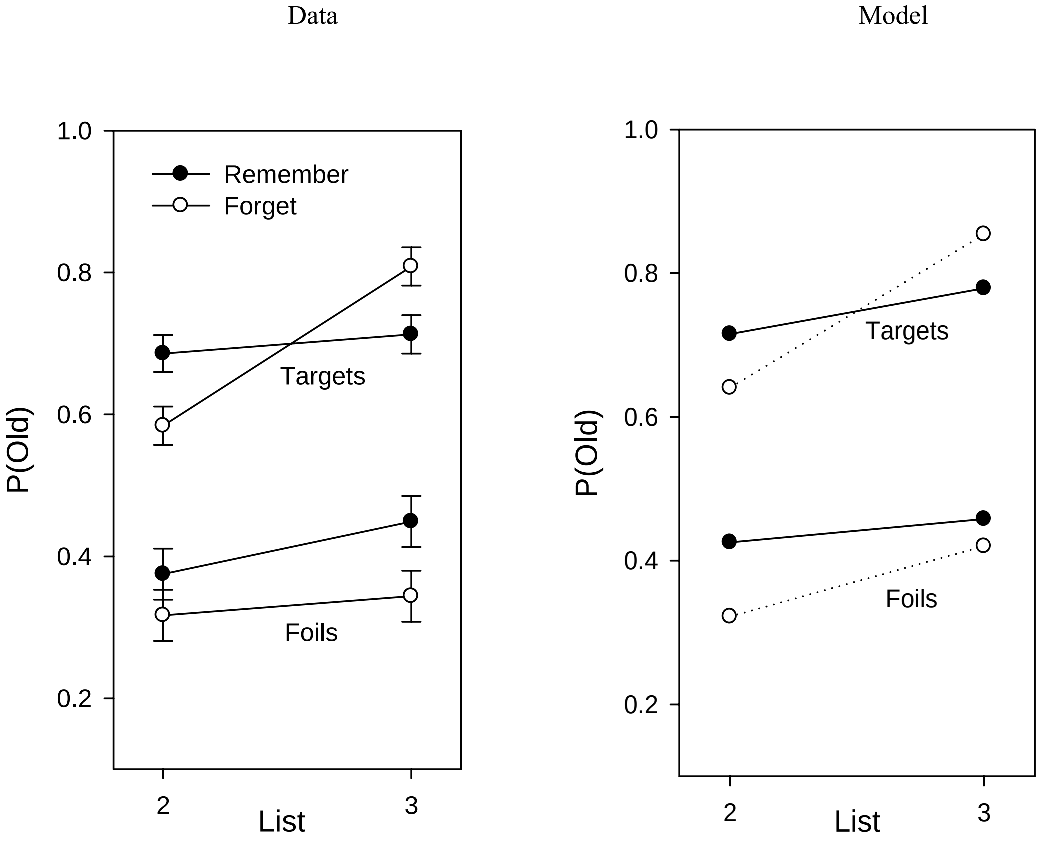

Figure 20. Model predictions for serial position data in exclusion recognition.

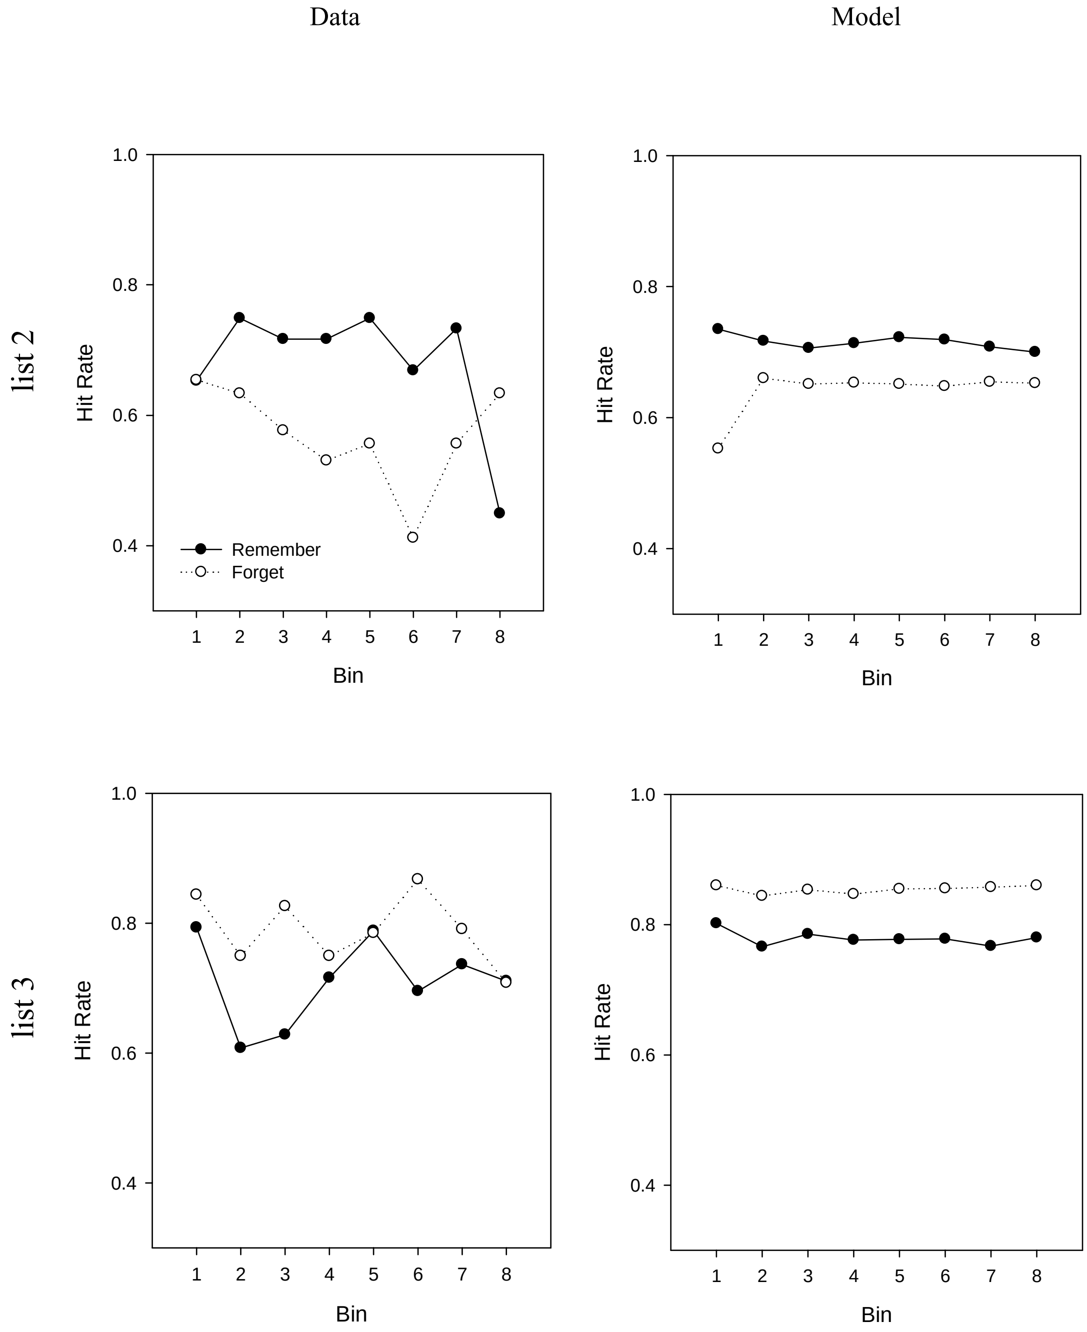

Modeling Additional Data

While not explicitly required of me, I also explored the ability of the model to account for the data from three additional experiments. Overall the model did well, especially when you note that the predictions derived for Experiment 6 simply fell out of the model derived for the earlier experiments. In that sense, Experiment 6 provided an *a priori* test of the model.

Recognition with shortened study time

The model parameters are identical to those in the previous recognition experiments aside from a reduced *t* value. Because study time is reduced by half, the *t* values for these two conditions were reduced by half. No other changes were made to the model. The new parameter values are: $t_1 = 3$; $t_0 = 1$; $t_{1(forget)} = 6$. Data and model predictions for Experiments 3 and 4 are presented in *Figure*s 21-25.

Delayed Free Recall

The model parameters are identical to those in the previous free recall experiment, except for an increased context change that occurs after the last list. In addition, because time has passed, it may be harder to recover the contents of a trace even after that trace is sampled. For this reason, recovery probabilities are also reduced by half. The new parameter value is: $\beta = .8$. Data and model predictions for Experiment 6 are presented in *Figure* 26.

Figure 21. Model predictions for inclusion recognition with 4 second study time.

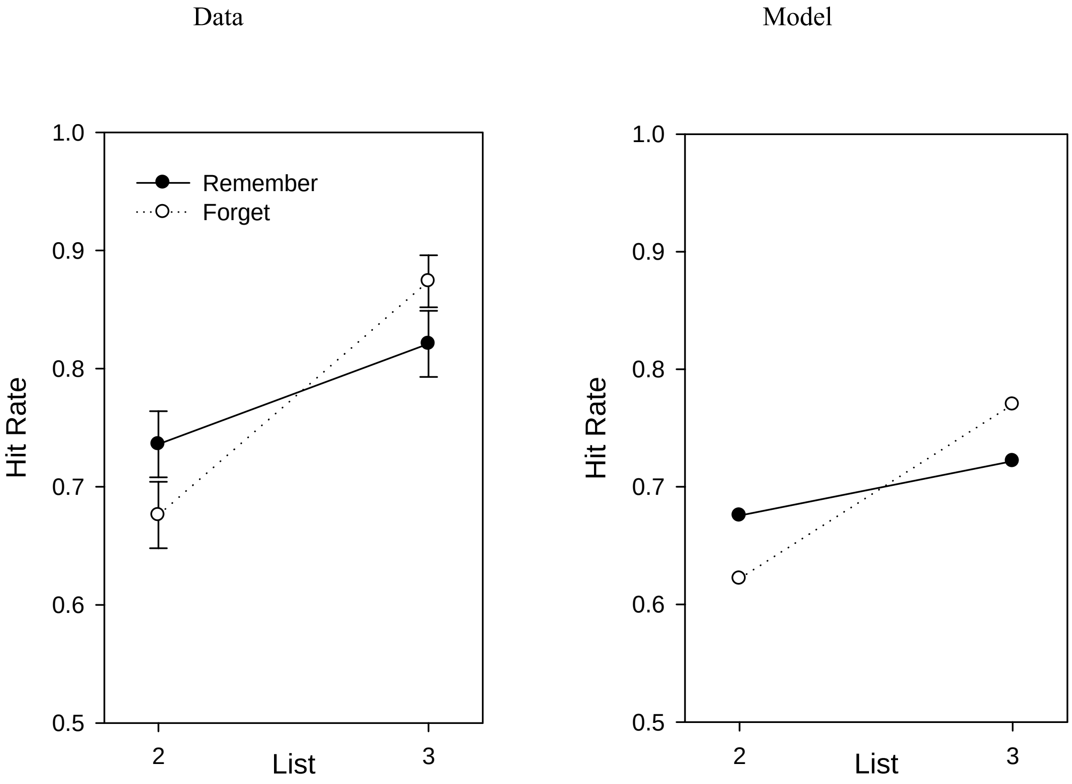

Figure 22. Model predictions for exclusion recognition with 4 second study time.

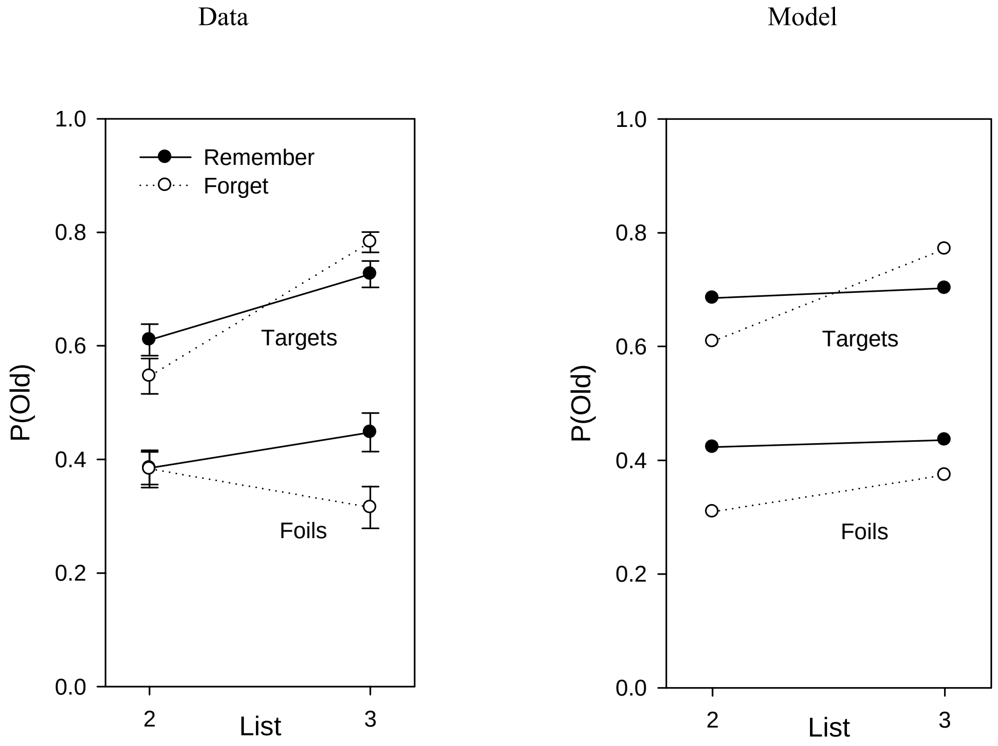

Figure 23. Model predictions for delayed free recall.

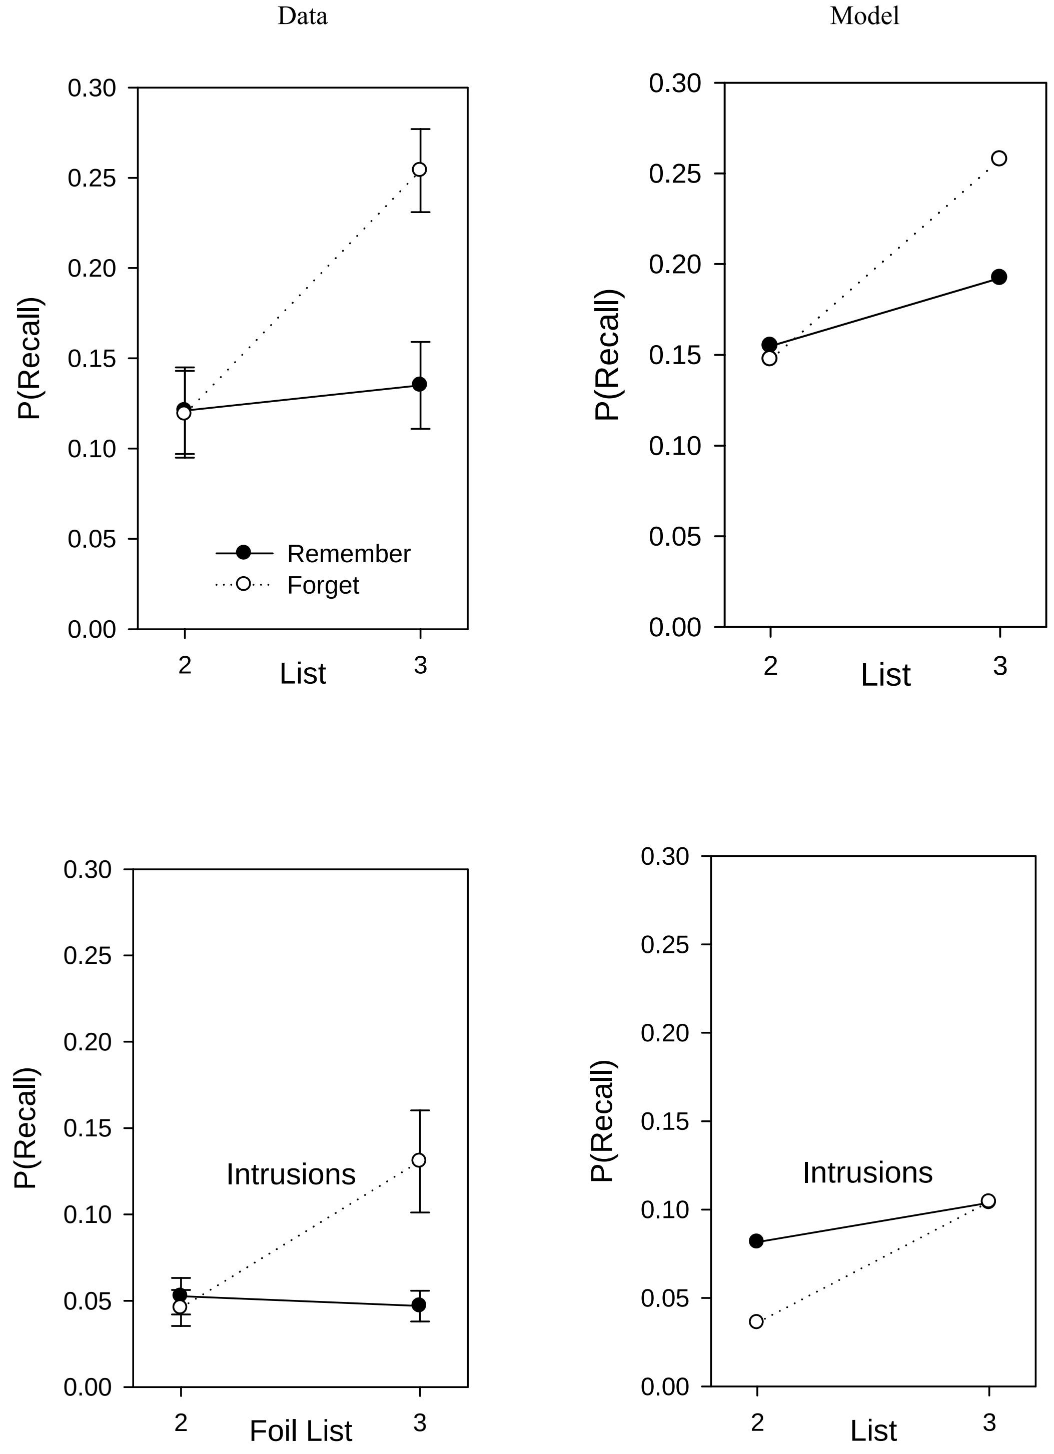

Figure 24. Model predictions for serial position data in delayed free recall.

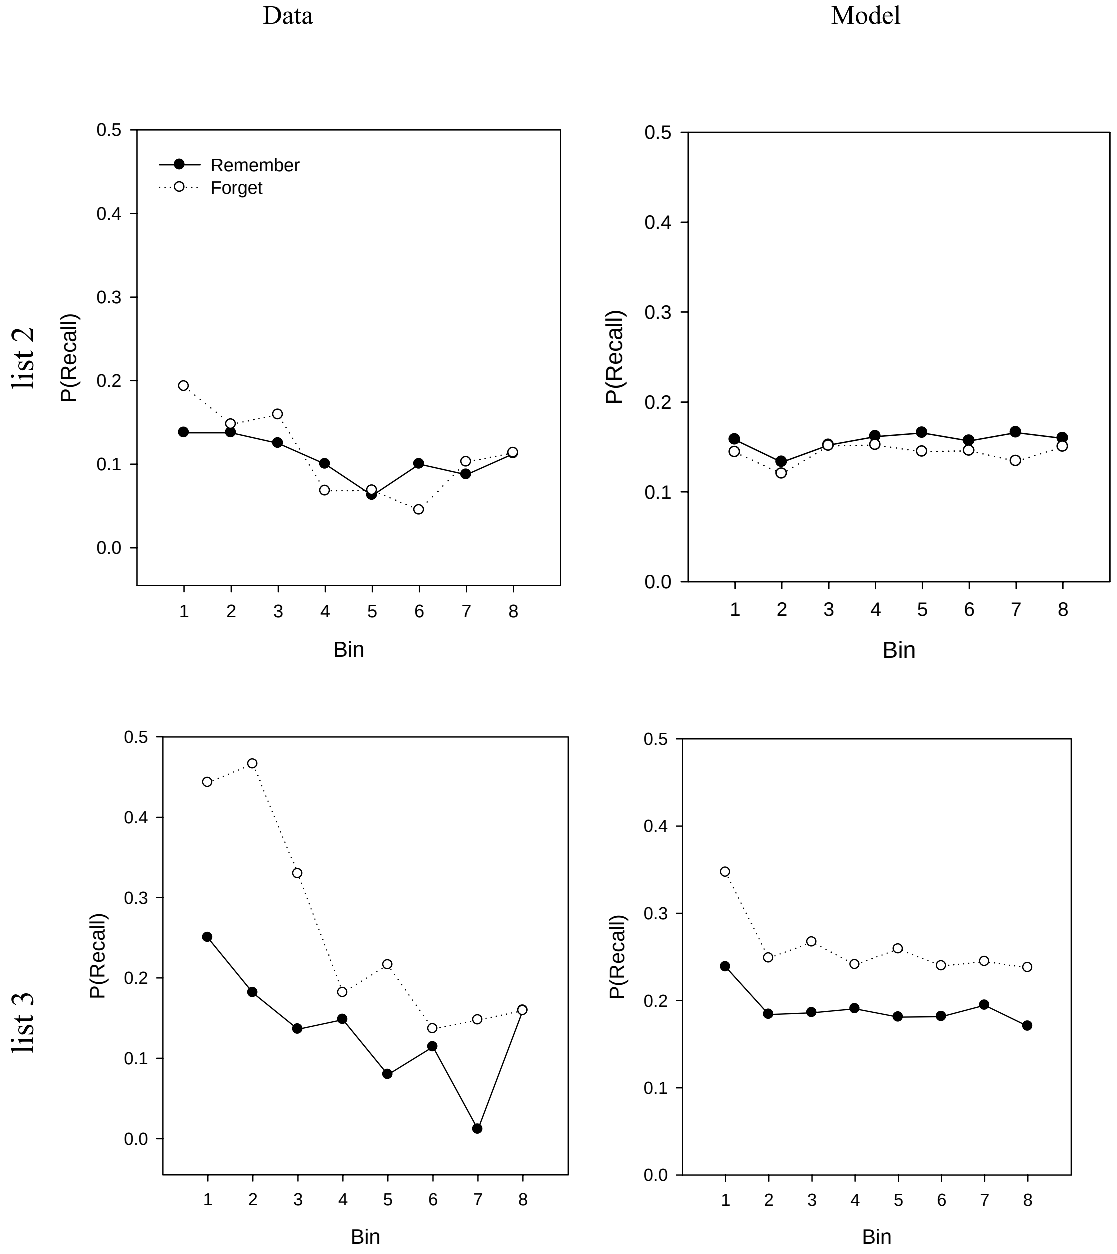

Note. For the sake of clarity, the 16 item list was compiled into 8 bins spanning two serial positions. For instance, bin *n* contains the data from serial positions *2n-1 and 2n.*

Figure 25. Model predictions for first item output position data in delayed free recall.

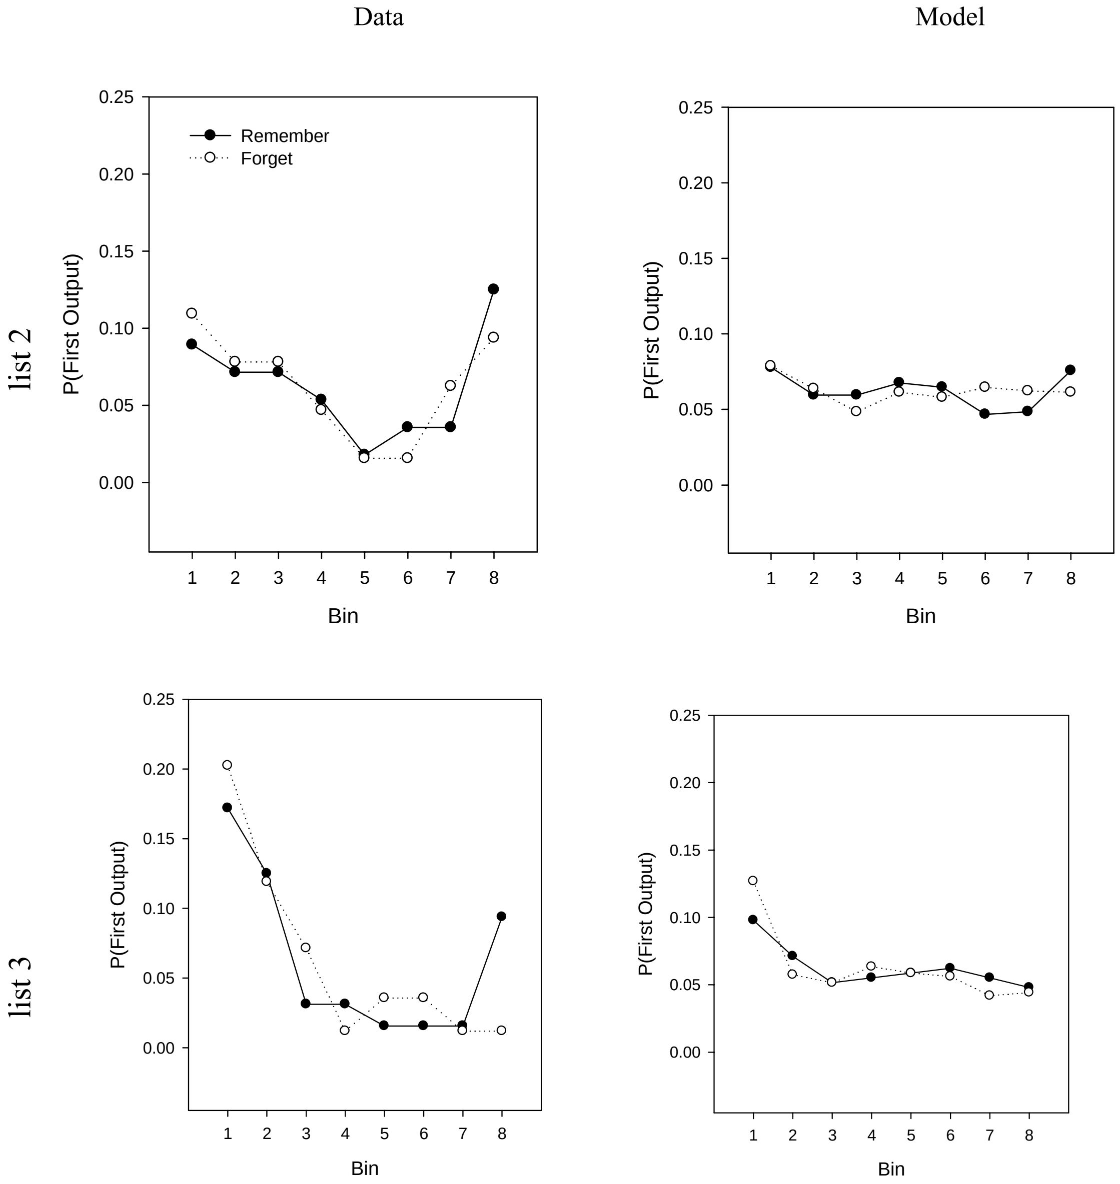

Note. For the sake of clarity, the 16 item list was compiled into bins. For first item output, *bin 1* represents the first item on the list (since this is where differences are seen) and all other serial positions are grouped by three.

Figure 26. Model predictions for conditional response probabilities from delayed free recall.

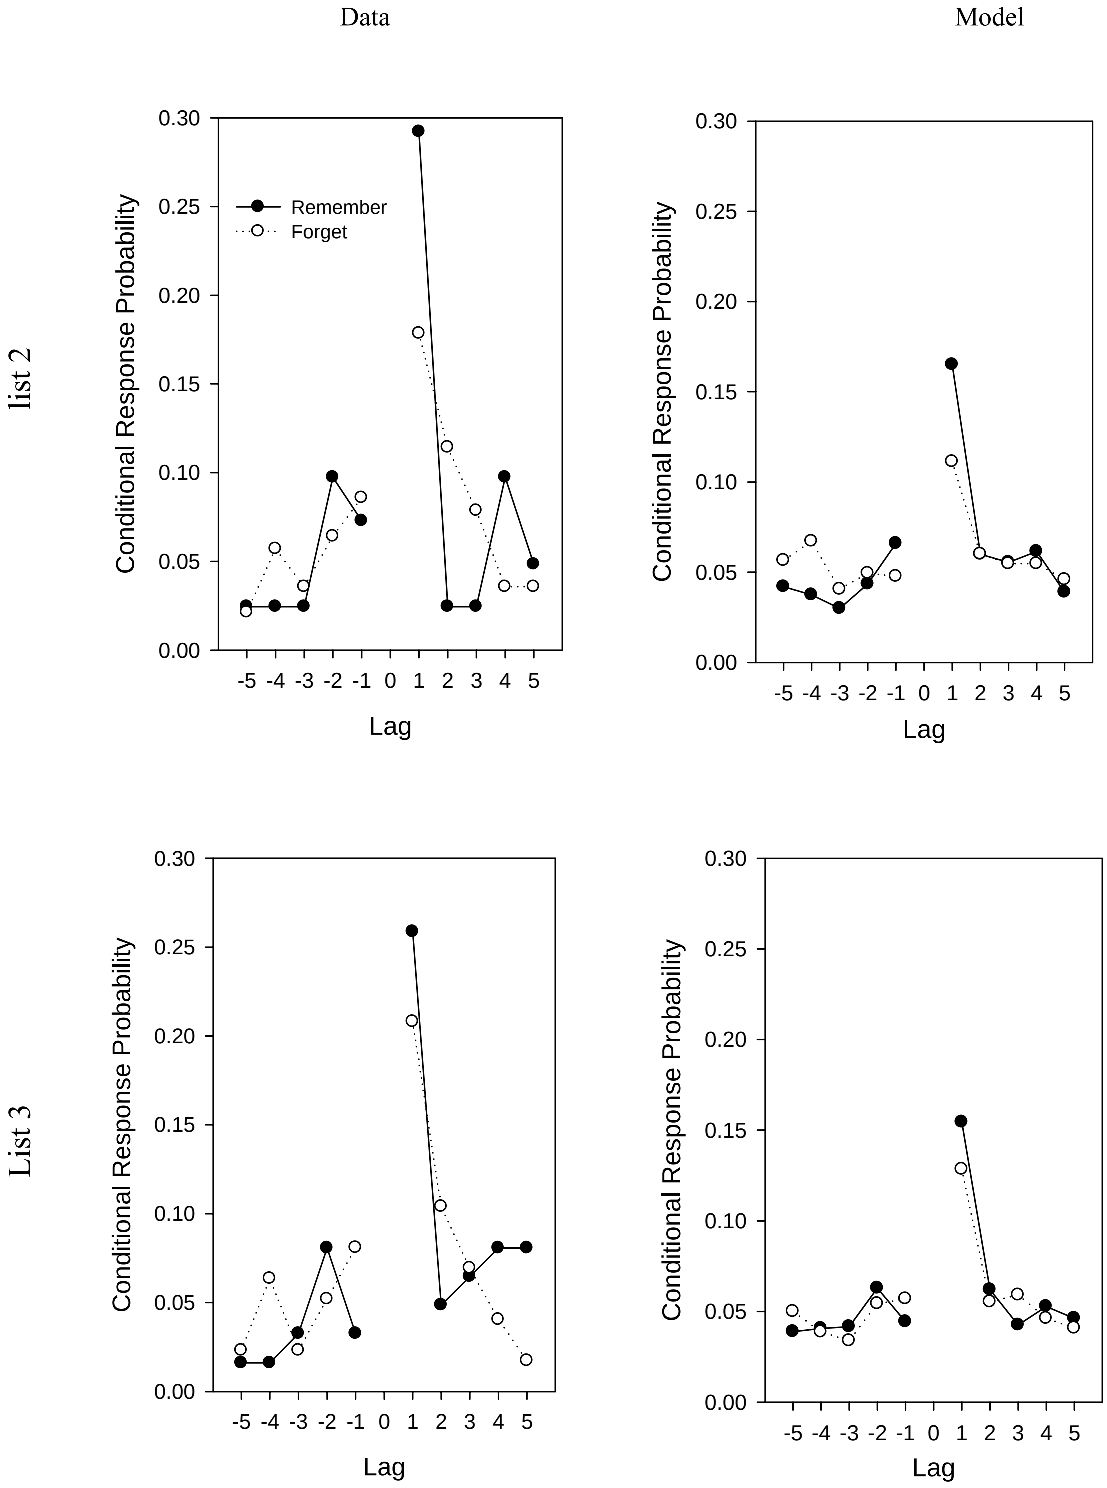

Chapter Four: General Discussion

The primary achievements of this project are both empirical and theoretical. By carefully considering the designs of prior experiments in the literature, I was able to develop hypotheses about the sources of many of the inconsistency observed there. When a set of carefully designed experiments were used to systematically explore directed forgetting, I observed reliable effects of directed forgetting for both recognition and recall. In addition, I was able to observe recency effects in free recall that to date had eluded prior investigations. With a coherent set of observations in place, I was then able to explore a wide variety of models with the REM framework to account for them. Here I will discuss the successes of the model already discussed and the failures of several models that lead to the current model.

The free recall model accurately predicted the patterns of data for both correct recall and intrusions. While the intrusion rates were higher overall in the model, the pattern of intrusions matched that of the data. The differences between remember and forget conditions in serial position, first item output, and CRP were sometimes smaller in the model compared to those in the data, however the overall patterns were consistent.

Aside from a few data points, model predictions were also consistent with the data for recognition inclusion and exclusion. For hit rates in inclusion, the model predicted both costs and benefits; the data showed costs, and the appearance of benefits that were not significant (perhaps due to a ceiling effect). Serial position predictions matched data for all lists except for L_2 in the forget condition, in which the model predicted a lower hit rate for the first item on the list compared to the rest of the items on the list; this pattern was not apparent in the data.

Both hit rates and false alarm rates seen in exclusion were predicted quite accurately by the model; the model matched the data both qualitatively and quantitatively. Results for model predictions in serial position were similar to those in inclusion – the model predicted the data well in all conditions except for L_2 in the forget condition, in which the model again predicted a lower hit rate for the first item on the list.

For inclusion and exclusion with reduced study time, the model was again accurate in predicting the data; the only difference predicted in the model that was not apparent in the data was a difference in the number of intrusions coming from L_2 during a test of L_3 between remember and forget conditions. The model predicted (as in free recall and exclusion with longer study time) an overall lower intrusion rate for the forget condition, but this was not apparent in the data.

Finally, the model predictions were also consistent with the data in delayed free recall. Data patterns were accurately predicted by the model for correct recall, serial position, first-item output, and CRP. Again, the degree of the effect may have been smaller, but was quantitatively predicted by the model. A slight difference was seen for intrusions; however the pattern of intrusion rates does not differ between the data in the model – intrusion rates for the remember condition are just lower overall in the data than in the model.

The model was successful in accounting for a significant amount of data given very few free parameters. Between the Remember and Forget conditions, only four parameters were changed, in order to account for 310 data points. In addition, by altering a few additional parameters, I was also able account for data from my recognition experiments with reduced study time, and my delayed free recall experiment.

There are a few critical aspects of the model contributing to the difference between remember and forget conditions. First, an increased context change between lists in the forget conditions creates less overlap in contextual features between L_2 and L_3. This contributes to both the costs and benefits of directed forgetting; the costs occur because the context at test shares less common features with L_2 and the benefits occur due to less competition from L_2 items when retrieving from L_3. The change in context that occurs with the forget instruction also leads to a difficulty in reinstating L_2 context features at the time of recall, also contributing to the costs. Finally, the decrease in overlap between the contexts of the two lists makes intrusion rates lower for the forget condition. Because contexts share less common features, the lists are more distinct, making participants better able to determine that a word came from an incorrect list.

In addition to creating a greater change in context, the forget instruction also has an effect on the rehearsal component of the model. After the forget instruction, the first item on L_3 is better encoded compared to the remember condition. This contributes not only to the benefits in the free recall experiment but also to the long lasting benefits in the delayed free recall experiment. Even after a delay in which context is again changed, the benefits persist, suggesting that L_3 items are better encoded. The serial position data from delayed free recall indicate that the initial items on L_3 in the forget condition are significantly more likely to be recalled than the initial items on L_3 in the remember condition, supporting the hypothesis that the first item is better encoded than other items.

My model of directed forgetting utilizes both context and rehearsal components. The free recall serial position and first-item output data discussed in this paper suggest that a combination of context change and rehearsal is contributing to the directed forgetting effect, and the current model provides the best predictions for these data. Further empirical work is necessary in order to properly evaluate these hypotheses. For example, experiments designed to eliminate rehearsal (either by using simultaneous tasks during encoding or incidentally encoded lists) within the current 3 list + distractor task design may help shed some light on the issue of rehearsal.

Future work on the model may also be needed in order to determine whether both context and rehearsal components are necessary. It seems unlikely that a pure rehearsal model could account for the data, given that a rehearsal model with no context change would not predict a recency effect, as there is no mechanism in a rehearsal model to produce recency. It may be that a combination context + rehearsal model that differs from the current one could account for the data; context may change over time, but without an increased context change after the forget instruction, and only a change in rehearsal contributing to the directed forgetting effect.

A context-only model may be a more reasonable model, however various context-only models were attempted, all of which were unable to predict the current patterns of the data. All of these models used the same basic process for directed forgetting as the current model – an increased context change between lists given the forget instruction. The first model developed was a context-only model which had context features that changed at different rates - some experimental context features stayed constant throughout all three lists, some changed more quickly so that there was less overlap in these features between lists, and some changed very rapidly so that they were different even for items in the same list. Thus, there were some features that were shared between lists and others that made the lists more distinct. The forget instruction increased the rate of change for all context features. Other variations in which the forget instruction only increased the rate of context change for certain features were also attempted, with similar results.

For retrieval, the rapidly changing features were not used in the sampling process (in an attempt increase sampling from L_2 or from the beginning of L_3). In a slightly different version of this model, these features were changed and random feature values were used instead. In either version, recency was present but I was unable to get higher sampling of L_2 in any case (meaning even when sampling was intended for L_2, L_3 items were still more likely to be sampled); at most, I could produce equal recall for L_2 and L_3. In addition, I was unable to get an advantage for items at the beginning of the list (primacy), a vital aspect of the model.

The second model used a vector that was divided into individual list components – one part of the vector was dedicated to L_1, one part was dedicated to L_2, and one part was dedicated to L_3. When encoding items from a given list, the sections of the vector that were dedicated to other lists would be left blank (with a few features being encoded erroneously). Two versions of this model were attempted – one where context changed between and within lists, and one where context changed between lists but not within. This model produced the costs, due to the context change making the test context less similar to that of L_2, but there was no mechanism to create the benefits, and thus they were not observed. This model also failed to produce primacy, as there was no mechanism to give the advantage to items in a specific position on the list.

A second version of this model was a combination context + rehearsal model. In this version, co-rehearsal of items in the buffer led to storage of item information from co-rehearsed items, and this information was used as part of the retrieval cue. Additionally, the first item on the list was given more rehearsals thus leading to better encoding. This eliminated the problem of no primacy (and also increased the probability of recalling successively studied items, which produced the CRP curves), but still did not produce the benefits of directed forgetting.

The next version of the model returned to the use of a single context vector, rather than one that was divided into list portions. During recall, the context of a given list was reinstated to use as a cue to recall from that list. This allowed for recall from a specific list but did not produce a serial-position curve. Another version of this model included special context features that were associated with only the first item on the list to be stored in lieu of co-rehearsed item features (since this item is alone in the buffer at the beginning of a study list). During recall, a special cue containing only these features was reinstated. This produced primacy, but again I was unable to get higher sampling of L_2 when it was the intended recall list.

Another manipulation of the model used the same “special” cue at the beginning of recall but then used a recovered item’s stored context as the context cue for sampling of the next item, a similar process to that used in Howard and Kahana’s (2002) Temporal Context Model (TCM). This model produced a primacy effect, however it also created recall that was too good – once an item from a given list was sampled, it was too easy to stay in that list. The final version of this model solved all of these problems by returning to the reinstated list cue (rather than a special cue for the first item on a list). Context information for the first item on a list was simply encoded better than that of other items on the list, and co-rehearsed item information was used as part of the cue after an item was recovered. This solved the previous problems but still did not allow us to get the first-item output patterns seen in the data. From this model, the current manipulations were made to produce a working model.

While none of the earlier models were able to fully account for all of the data, a combination of one of these models combined with the current assumptions may be necessary. A variation of the current model in which context changes within a list (in addition to between lists) may be better able to account for the data. While this manipulation was present in earlier models that were not successful in predicting the directed forgetting data, adding this assumption into the current model may provide a better account of the current data and may also better predict future data.

If I implement the assumption that context changes within a list, it will also be important to consider the way that context changes within a list. Traditionally, models of context change within a list (Mensink & Raaijmakers, 1989) assume that context fluctuates randomly throughout the list. In my model, this would be represented by the same type of change that occurs between lists occurring within lists at a slower rate. Howard and Kahana’s (2002) Temporal Context Model (TCM) assumes that rather than context changing randomly, context drifts based on contextual states elicited by studied items.

My model was quite successful in accounting for a variety of data given very few free parameters. It will be useful to examine other ways in which this model can be manipulated to better fit the data (for example by adding in a within-list context change similar to that in TCM). In addition to explaining directed forgetting data, a context-change model such as this could be used to explain other context-dependent phenomena.

Given the success of the current model, it will also be possible to generate predictions concerning the effects of directed forgetting on other specific memory tasks, which of course can be empirically tested. I am particularly interested in extending the model to explain the effect of directed forgetting on memory performance measured using the item-method. As previously discussed, there is no reason to believe that the item-method would necessitate a context change component to the model, and I predict a directed forgetting effect using only a rehearsal manipulation in the item-method. A rehearsal model for the item method would eliminate all assumptions about context-change and instead utilize a manipulation of number of rehearsals each item receives. This effect would be achieved by changing the value of the *t* parameter between to-be-remembered and to-be-forgotten words.

References

Atkinson, R.C., G.H. Bower and E.J. Crothers, *An Introduction to Mathematical Learning Theory*, New York: Wiley, 1965.

Atkinson, R.C. & Shiffrin, R.M. (1968). Human memory: A proposed system and its control processes. In Spence, K.W. & Spence, J.T. (Eds.), *The Psychology of Learning and Motivation (pp 89-195),* New York: Academic Press.

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.

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

Bjork, R.A., LaBerge, D., & Legrand, R. (1968). The modification of short-term memory through instructions to forget. *Psychonomic Science, 10,* 55-56.

Bjork, R.A., & Whitten, W.B. (1974). Recency-sensitive retrieval processes in long-term free recall. *Cognitive Psychology, 6*, 173-189.

Block, R. A. (1971). Effects of instructions to forget in short-term memory. *Journal of Experimental Psychology, 89 ,* 1–9.

Brown, J. (1958). Some tests of the decay theory of immediate memory. *The Quarterly Journal of Experimental Psychology*, 10, 12-21.

Ebbinghaus, H. (1885). *Memory: A contribution to experimental psychology*. New York: Teachers College, Columbia University.

Eich, J. E., Weingartner, H., Stillman, R. C., & Gillin, J C. (1975). State-dependent accessibility of retrieval cues in the retention of a categorized list. *Journal of Verbal Learning & Verbal Behavior, 14*(4), 408-417.

Elmes, F. J., Adams, C., & Roediger, H. L. (1970). Cued forgetting in short-term memory: Response selection. *Journal of Experimental Psychology, 86,* 103–107.

Epstein, W. (1969). Poststimulus output specification and differential retrieval from short-term memory. *Journal of Experimental Psychology, 82,* 168-174.

Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*, 62, 145–154.

Francis, W. N., & Kucera, H. (1982). *Frequency analysis of English usage: Lexicon and grammar.* Boston: Houghton Mifflin.

Geiselman, R. E., Bjork, R. A., & Fishman, D. (1983). Disrupted retrieval in directed forgetting: A link with posthypnotic amnesia. *Journal of Experimental Psychology: General, 112,* 58–72.

Glanzer, M. & Cunitz, A. R. (1966). Two storage mechanisms in free recall. *Journal of Verbal Learning & Verbal Behavior, 5*(4), 351-360.

Godden, D., & Baddeley, A. D. (1975). Context-dependent memory in two natural environments: On land and under water. *British Journal of Psychology, 66,* 325–331.

Godden, D., & Baddeley, A. D. (1980). When does context influence recognition memory? *British Journal of Psychology, 71,* 99–104.

Goodwin, D. W., Powell, B., Bremer, D., Hoine, H., & Stern, J. (1969). Alcohol and recall: State dependent effects in man. *Science, 163,* 1358–1360.

Jacoby, L. L (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30*(5), 513-541.

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

Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition*, 24, 103–109.

Howard, M.W. & Kahana, M.J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology, 46*, 269-299.

Kimball, D.R., Smith, T.A., & Kahana, M.J. (2007). The fSAM model of false recall. *Psychological Review, 114*(4) 954-993.

Macht, M. L., Spear, N. E., & Levis, D. J. (1977). State-dependent retention in humans induced by alterations in affective state. *Bulletin of the Psychonomic Society, 10,* 415–418.

MacLeod, C.M. (1975). Long-term recognition and recall following directed forgetting. *Journal of Experimental Psychology: Human Learning and Memory*, *1*, 271–279.

MacLeod, C.M. (1989). Directed forgetting affects both direct and indirect tests of memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 13–21.

MacLeod, C. M. (1998). Directed forgetting. In J. M. Golding & C. M. MacLeod (Eds.), *Intentional forgetting: Interdisciplinary approaches* (pp. 1–57). Mahwah, NJ: Erlbaum.

MacLeod, C. M. (1999). The item and list methods of directed forgetting: Test differences and the role of demand characteristics. *Psychonomic Bulletin & Review, 6*(1), 123-129.

Malmberg, K. J., Holden, J. E., Shiffrin, R. M. (2004). Modeling the Effects of Repetitions, Similarity, and Normative Word Frequency on Old-New Recognition and Judgments of Frequency. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30*(2), 319-331.

Malmberg, K. J. & Shiffrin, R. M. (2005). The 'One-Shot' Hypothesis for Context Storage. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*(2), 322-336.

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 for contextual fluctuation. *Journal of Mathematical Psychology*, 33, 172–186.

Murnane, K., Phelps, M. P., & Malmberg, K. (1999). Context-dependent recognition memory: The ICE theory. *Journal of Experimental Psychology: General, 128,* 403–415.

Peterson, L. R., & Peterson, M. J. (1959). Short term retention of individual verbal items. *Journal of Experimental Psychology*, 58, 193-198.

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

Rundus, D. (1971). Analysis of rehearsal processes in free recall. *Journal of Experimental Psychology, 89*(1), 63-77.

Sahakyan, L. (2004). The destructive effects of the “forget” instructions. *Psychonomic Bulletin & Review, 11,* 555–559.

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*(4), 789-801.

Sahakyan, L., & Kelley, C. M. (2002). A contextual change account of the directed forgetting effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28,* 1064–1072.

Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin & Review, 4,* 145–166.

Shiffrin, R. M., & Steyvers, M. (1998). The effectiveness of retrieval from memory. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition* (pp. 73–95). London: Oxford University Press.

Slamecka, N.J. & McElree, B. (1983). Normal forgetting of verbal lists as a function of their degree of learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9*(3), 384-397.

Smith, S. M. (1979). Remembering in and out of context. *Journal of Experimental Psychology: Human Learning and Memory, 5,* 460–471.

Winograd, E. (1968). List differentiation, recall, and category similarity. *Journal of Experimental Psychology*, *78*(3), 510-515.

Wixted, J. T. & Stretch, V. (2004). In defense of the signal-detection interpretation of Remember/Know judgments. *Psychonomic Bulletin & Review, 11*, 616-641.

Woodward, A.E., Jr., & Bjork, R.A. (1971). Forgetting and remembering in free recall: Intentional and unintentional. *Journal of Experimental Psychology, 89,* 109-116.

Xu, J. & Malmberg, K. J.(in press). Modeling the effects of verbal- and non-verbal pair strength on associative recognition, *Memory & Cognition.*