Journal of Memory and Language **43,** 409–430 (2000)
doi:10.1006/jmla.2000.2706, available online at http://www.idealibrary.com on IDEAL®

The Disruption and Dissolution of Directed Forgetting: Inhibitory Control of Memory

Martin A. Conway, Kay Harries, Jan Noyes, Mihaly Racsma'ny, and Clive R. Frankish

Department of Experimental Psychology, Centre for Learning and Memory, University of Bristol, Bristol England

In a series of directed-forgetting (DF) experiments it was found that inhibition of a to-be-forgotten (TBF) list could be disrupted by a secondary task and completely abolished by a concurrent memory load during second to-be-remembered (TBR) list learning. Similarly, inhibition was found to be wholly abolished when the TBF and TBR list were strongly associated but not when weakly associated. These findings suggest that inhibition in the DF procedure depends on how powerfully the second TBR list competes in memory with the representation of the TBF list. When the representation of the TBR list is impoverished or when it is too similar to the TBF list then competition is weak and inhibition is as a consequence weak or does not occur at all.

Key Words: long-term memory; inhibition; directed forgetting; encoding; output interference; autobiographical memory; free recall; recognition; control processes; forgetting.

A fundamental problem facing any memory system is how to select individual items from a range of related items, all of which are (equally) available for recall. For instance, recalling the comments of person A from yesterday's meeting rather than those of B and C, recalling fact X rather than the related facts Y and Z, or, more mundanely, simply recalling where one parked one's car today as opposed to yesterday or the day before (Bjork, 1989) all require resolution of competing responses in memory. This problem may be especially acute for very recently encoded knowledge, much of which may be in a highly accessible state for some limited period after encoding. Short-term retention, over a period of hours or possibly longer, of a very detailed record of experience may occur because it is not possible at the moment of encoding to fully determine what should be retained and what can be forgotten and, therefore, an "overly" detailed record is initially represented in memory. The idea being that over time personally relevant information will receive more frequent rehearsal, i.e., will be recalled or accessed more frequently than less personally relevant details and, as a result, will enter a more stable and durable state. In contrast, other less frequently accessed details of a memory will become, over time, progressively more difficult to access. One reason for this is that each time a single or set of details are accessed associated details are inhibited (Anderson, Bjork, & Bjork, 1994; Anderson & Spellman, 1995; Radvansky, 1999). Nevertheless, it seems that inhibition, at least in the short-term, can be overcome by a sufficiently specific cue, i.e., as in a recognition test, and this allows the upgrading of accessibility of details which subsequently turn out to have a significance not apparent at encoding. In this way access and the inhibition it causes "shape" a memory into a pattern of highly accessible details and strongly inhibited, inaccessible, details. Presumably, the latter eventually reach some asymptote of inhibition at which they become permanently inaccessible, even to specific cues, and so to all effects are forgotten.

Inhibition may then shape detailed memories over subsequent episodes of recall. It would,

This research was supported by an ESRC training award, R00429634177, to Kay Harries and by the Department of Experimental Psychology, University of Bristol. The authors thank both institutions for their support. We also thank Bob Bjork for discussion of, and helpful revision to, this article.

Address correspondence and reprint requests to Martin A. Conway, Department of Experimental Psychology, University of Bristol, 8 Woodland Road, Bristol, BS8 1TN England. Fax: +117 928 8588. E-mail: M.A.Conway@bristol.ac.uk.

however, be incorrect to assume that for any experience all or most features are retained, even in the short-term, in identical states of high and equal accessibility. Instead, our personal goals, beliefs, motives, and comprehension of the world ensure that knowledge of any individual experience is initially stored according to how it relates to the self. Indeed, the degree of attention given to specific aspects of an experience, the predictability of an event (its "script-likeness"), and other purely endogenous factors, i.e., time of day, will all act to determine accessibility of different details of an experience. Thus, even though a large amount of knowledge of an experience may initially be retained, this will be, at encoding, already in a pattern of accessibility in which some features are high in terms of accessibility, whereas others are low or barely accessible at all. In general, it seems to us that the aim of encoding is to create representations in memory that prioritize access to information relevant to the self (cf. Conway, 1996; Conway & Pleydell-Pearce, 2000) and minimize access to irrelevant or redundant information while at the same time preserving the availability of many details so that those which only become self-relevant later can still enter a state of high accessibility. The experiments reported in this paper address these encoding issues and argue that inhibitory processes operating during and shortly after encoding act to shape patterns of recently acquired knowledge in such a way as to give memories their (initial) form. We propose that this "shaping" of memories by inhibitory processes is a nonconscious, automatic process that creates memories in which some knowledge is more accessible than other knowledge when a memory is subsequently recalled. Although we conceive of this process as occurring outside consciousness it may, at least on some occasions, require sustained conscious attention in order to operate effectively. Indeed, in all the experiments reported below we used a procedure known as *directed forgetting* (DF) in which people were overtly instructed to self-initiate the forgetting of recently acquired materials. It is to this procedure that we turn next.

INHIBITION IN DIRECTED FORGETTING

In the DF procedure that uses lists of words participants are simply informed that they will see a list of words which they will later have to remember. After presentation of half the list one group of participants are instructed to forget the words they have just learned (the F group). In contrast, another group are instructed to keep remembering the first lists of words (the R group). Both groups then learn a second list which they are informed they will have to later recall. At test participants are required to free recall all previously studied items. Under these conditions the F group's recall of list 1 is at a low level, typically in the range 30 to 40%, whereas, and in striking contrast, recall of list 1 by the R group is unimpaired and at a high level that is typically between 60 and 80%, depending on the delay between study and test (Bjork, 1989; MacLeod, 1998). Exactly how this self-initiated forgetting occurs is unknown, although current theory favors an account in terms of inhibition (Bjork, 1989; Bjork, Bjork, & Anderson, 1998; MacLeod, 1998). By this view, when, for example, a recently acquired list of words is followed by an F cue then this triggers inhibitory mechanisms which act to reduce the accessibility of the list's representation in long-term memory so that few items are accessible in free recall. This process of *retrieval inhibition* (Bjork, 1989; Bjork, Bjork, & Anderson, 1998) lowers the accessibility of otherwise normally acquired items while leaving their availability unaffected (Tulving & Pearlstone, 1966). As a consequence, retrieval inhibition can be abolished if specific cues—cues more specific than those present in free recall—are presented at test, i.e., the actual list items themselves, as in a recognition test. But note, that mere exposure to the inhibited list items does not overcome the inhibition; rather, for this to occur the items must be represented in the context of intentional remembering (Bjork & Bjork, 1996). An F cue, then, presented after acquisition of a list in a DF experiment triggers inhibitory processes that lower the accessibility of the newly acquired list items, although these remain available and can be retrieved when specific cues enter the retrieval process. The result is a representation in which (episodically) associated information is differentially accessible although all (retained) information is, at least initially, equally available. Thus, the DF procedure provides a way in which the process of establishing differential memory content accessibility can be studied in the laboratory.

There are a range of conditions that influence the occurrence and magnitude of DF (see MacLeod, 1998, Table 1.2; and Bjork et al., 1998, for reviews) and of these one of the most important is that, in list-learning DF experiments, the F cue *must* be followed by the learning of a new list. Without second-list learning there is no reliable impairment in recall of the first list (Bjork, 1989) and, thus, learning the second list is critical to the DF effect. Why second-list learning is so crucial is not known but from participants's self reports it seems that learning the second list provides an opportunity to focus attention on items other than the TBF items. It may be that this focusing of attention on the second TBR list is critical to the DF effect because it triggers inhibition of other, already acquired, TBF items which are competing for attention (Anderson, Bjork, & Bjork, 1994; Anderson & McCulloch, 1999; Anderson & Spellman, 1995; Bjork, 1989; Bjork et al., 1998; Tipper, 1985; Tipper & Driver, 1988). Anderson et al. (1994) demonstrated the inhibitory effect of competitor items in the retrieval practice procedure in which prerecall rehearsal of an item from a recently studied list causes inhibition of related list items in a later recall test, e.g., rehearsal of Orange (cued by the category name plus word fragment, "Fruit-Or . . . ?") from the earlier acquired list "Apple, Orange, Banana" causes inhibition of the competitor items Apple and Banana. Anderson and Spellman (1995) extended these findings and showed how spreading inhibition could extend beyond category boundaries to other associates, e.g., prerecall rehearsal of the pair Green–Emerald from an earlier acquired list that included Green–Emerald, Green–Lettuce, Soup–Mushroom causes inhibition of both Lettuce *and* Mushroom on a later recall test. This finding of spreading inhibition, or second-order inhibition (Anderson & Spellman, 1995), is important because it rules out a range of alternative explanations of these effects such as accounts in terms of occlusion, interference, and so forth (see Anderson & Spellman, 1995; Anderson & Bjork, 1994; and Bjork et al., 1998, for further discussion). The effects of inhibition in the retrieval practice procedure are, then, extensive and arise because of competition between already-encoded list items with the to-be-rehearsed retrieval practice items. Our suggestion is that a similar form of competition occurs between the TBF and TBR lists in the DF procedure and it is this which triggers (automatic) inhibitory processes as attention is intentionally and selectively focused on the second list.

DISRUPTION OF INHIBITION

A counterintuitive prediction arising from the above reasoning is that if attentional resources are stressed or overloaded during the learning of the second list then it may not be possible to focus attention on the TBR items and so trigger inhibition of the TBF competitors. A consequence of this would be an unexpected rise in the recall of the F list, which would not have been effectively inhibited. Interestingly, this was exactly the finding of Macrae, Bodenhausen, Milne, and Ford (1997), who increased the processing demands in second-list learning by requiring F-group participants to keep a running total of all the vowels in the words on the second TBR list. Unsurprisingly, memory for list 2 was markedly impaired with a mean probability of recall of .22 compared to a mean of .49 for the F list. In contrast, mean recall of the F list in the condition where there was no secondary task on list 2 was significantly poorer, at .37. It should be noted that the Macrae et al. DF procedure and materials were slightly different from those typically used and this was because the primary purpose of their study was to examine the processing of stereotypes in memory. Nevertheless, these findings show that increasing processing demands during second-list learning can *disrupt* DF and perhaps this occurs by attenuating attentional focusing during list 2 learning, thereby restricting the strength of inhibition of competitors (F-list items).

Another way in which to disrupt or even abolish inhibition would be to reduce the extent of competition between lists or items. Anderson and McCulloch (1999) found that by increasing the integration of items in the initial list the inhibitory effects of retrieval practice could be reliably reduced. Thus, simply having participants interrelate items at study, i.e., note that in the list animal–dog, animal–cat, the exemplars (dog and cat) are both household pets, significantly reduces the inhibitory effects cause by prerecall retrieval practice. Anderson and McCulloch (1999) were able to demonstrate that this effect arose solely because of this "lateral" integration of items during initial learning. A related set of findings has been reported by Radvansky (1999) in a series of studies that used the fan effect (Anderson, 1974). The fan effect refers to increasingly slower recognition times to items that have increasing numbers of associates. For instance, the time taken to verify the statement "The hippie is in the park" increases with the number of other facts known about the "hippie" and/or the "park." If, however, the items are integrated, e.g., the glass door is in the hotel, the welcome mat is in the hotel, the counter is in the hotel, then the fan effect, the slowing of recognition speed with increasing number of related facts, is not observed. In contrast, if the items are unintegrated, e.g., the welcome mat is in the hotel, the welcome mat is in the office, the welcome mat is on the porch, then reliable fan effects are observed. Radvansky (1999) argues that these results show the inhibition of unintegrated but competing items, competing presumably because they were learned at the same time, and a lack of inhibition for integrated items which do not compete with each other. Finally, Golding, Long, and MacLeod (1994), in an item-by-item DF experiment, found that when an F item could be integrated with an upcoming R item, e.g., High (F), Way (R), then DF on a later recall test was reliably reduced. Taken together these findings lend strong support to the view that materials which can be integrated at encoding have some immunity to inhibition caused by subsequent instructions or memory access.

It seems likely then that the distribution of attentional resources at encoding could be an important determinant of the extent of inhibition in the DF lists procedure. Similarly, the degree of integration of F materials may have a significant impact on the degree of competition between items and, therefore, upon the extent of inhibition. Indeed, inhibition and integration may be complimentary in that when integration at encoding is extensive, then inhibition is either limited or not present at all. Whereas when integration is limited or not possible, then inhibition may be strong and widespread—we return to these points later. In the experiments which follow we explore how attenuating attentional resources in second-list learning can, paradoxically, disrupt and even abolish inhibition of the F list. In particular, we replicate the Macrae et al. study using a more standard DF procedure and we then systematically explore disruption of inhibition. We also consider the wider implications of these findings in, for example, the formation of everyday autobiographical memories and and in the formation of highly vivid memories of traumatic experiences. Later experiments investigate how integration, too, can overcome inhibition of the F list and we extend the Golding et al. item-by-item study to disrupted inhibition of interrelated lists. The findings from these final studies are also extended to a discussion of the formation of autobiographical memories in which integration is a critical, but rarely considered, encoding and consolidation process.

GENERAL METHOD

The experiments used a standard procedure as described in this section. Where an experiment departs from this procedure exact changes are specified with the description of the experiment.

Design

A mixed design was used with one between-subjects factor and one within. The between-subjects factor, group, had two levels: the F group, who were given a between-lists F cue, and the R group, who received a between-lists R cue. The within-subjects factor was the repeated measure, lists, with two levels: list 1 (studied first) and list 2 (studied second). Order of presentation of lists was counterbalanced across participants and words were randomly allocated to lists and to experiments. Each study phase list contained 10 items. The list items were randomly drawn from a pool of 40 words of moderate to high frequency, naming common objects and locations.

Participants

There were 24 participants in each experiment and they were randomly assigned to F and R groups, each of which had 12 members. The participants were all student volunteers and were paid £2. Their age ranged from 16 to 45 with a mean of 18.4 years (only 3 participants were over 40 years of age). No attempt was made to select for gender and the numbers of men and women assigned to groups and experiments was unsystematic, although by chance there were approximately two more women than men in each experiment.

Procedure

Participants were informed that they were taking part in a memory experiment and would study two short lists of words which they should commit to memory. They were tested in small groups (either F or R but not mixed) and each person was seated at a separate computer console. The computers were sufficiently apart and angled so that participants could view only their own screen on which the TBR words were displayed for 2 s with a 2-s interitem interval. Participants initiated display of the first list when instructed by the experimenter to press the space bar on the computer keyboard. The items were then displayed automatically and after the 10th word the message “End of List 1” was displayed and the display ceased. At this point participants in the F group were informed that the list they had just heard “was in fact a practice list to familiarise you with the presentation rate and type of words. You should now put these words out of mind, try to forget them and not let them interfere with learning the experimental list which will be presented now.” The experimenter then instructed participants to press the space bar to initiate presentation of the second list. The same procedure was followed for the R group, who, rather than receiving an F cue were informed: “That is the end of the words on list one. You must try to keep those in mind as you learn the second list which will be presented now” and participants then pressed the space bar to initiate list 2 presentation. As soon as presentation of the second list was complete participants were handed several sheets of arithmetic problems and ask to solve as many as possible in the next 5 min. Following this participants were given a single sheet of paper and asked to recall any words they could from both lists and to start at the top of the page and write each word recalled under the previous word.

After completing free recall, which took on average 8 min, participants took a recognition test. The recognition test consisted of the 20 items from lists 1 and 2 and the remaining 20 unused words from the stimuli pool. Old and new items were intermixed in an unsystematic order and order of presentation was varied unsystematically across participants. Each word was listed along with (i) the words “OLD”–“NEW” (please circle), (ii) a memory-awareness scale (taken from Conway, Gardiner, Perfect, Anderson, & Cohen, 1997), and (iii) a separate 5-point confidence rating scale. For the memory-awareness judgments participants were given the following instructions:

> Indicate what sort of memory awareness you had when making an Old judgment. For instance you might have judged a word as old because you remembered it’s occurrence in one of the earlier lists. Perhaps, you recalled what you thought when you read the word earlier, perhaps you remember it preceded or was followed by some other word, you may even have an image of the word on the screen. All these forms of awareness are part of the Remember state. Conversely, you may experience none of these attributes of remembering but nonetheless have a powerful feeling of knowing that a word was seen in one of the earlier lists. This is the Just Know state and if you do not have an experience of remembering when you make your Old/New judgment then you may find you Just Know that a word was seen before. Finally, you may judge a word to be Old not because you Remember or Just Know it but, rather, because it feels familiar. When you make an Old judgment of a word you must then indicate whether the basis for your judgment was that you remembered, (just) knew, or felt the word to be familiar.

The experimenter checked verbally with the participants that they understood these response categories. On the confidence rating scale participants rated how confident they were that their Old/New judgments were correct and 1 = *Guess*, 2 = *Slight Confidence*, 3 = *Moderate Confidence*, 4 = *Strong Confidence*, and 5 = *Certain.*

EXPERIMENT 1

In the first experiment we simply carried out a "standard" DF experiment following the general method specified above. The purpose was to provide a baseline against which to compare later experiments. The DF effect is characterized by the following pattern of differences: in the F group list 1 recall is reliably poorer than recall of list 2, F group recall of list 1 is reliably poorer than R group recall of list 1, and F group recall of list 2 is higher than R group recall of the same list [presumably due to reduced proactive inference (PI) from the inhibited list 1 in the F group]. In the present experiments various additional manipulations occur during list 2 learning which sharply reduce overall recall of these items. Because of this we do not expect to find consistent across-experiment differences in which F-group list 1 recall is always significantly poorer than list 2 performance. The reduction in PI due to F-group list 1 inhibition is not always observed (MacLeod, 1998) and the conditions under which it does and does not occur are not known (Bjork et al., 1998). Accordingly we do not expect to observe this difference in all experiments. However, we do expect to observe a reliable and consistent difference between F-group list 1 performance and R-group list 1 performance in which the latter regularly and significantly gives rise to higher recall rates than the former: it is this difference that we take as the defining feature of the DF effect. We predict that none of these inhibitory effects will be present in recognition which usually overcomes the inhibition of list 1 items (MacLeod, 1998). Finally, an intriguing possibility is that although the F group and R group both perform at the same high level on the recognition test the F groups correct recognition of list 1 items will be associated with feelings of familiarity, whereas the R group's recognition of the same items will be associated with recollective experience (Gardiner, Gawlik, & Richardson-Klavehn, 1994). For expository purposes we report the free-recall data for each experiment first, followed by analyses of output order, and in General Discussion we report the recognition data for all experiments.

TABLE 1

Free Recall of Lists 1 and 2: Mean Probabilities for Experiment 1

	List 1	List 2
Experiment 1: "Standard" DF procedure		
R group	0.67	0.49
F group	0.32	0.51
Experiment 1a (5-item lists, $N = 11$)		
R group	0.69	0.68
F group	0.44	0.71
Experiment 1b (5-item lists, $N = 14$)		
R group	0.64	0.49
F group	0.36	0.74
Experiment 1c (20-item lists, $N = 24$)		
R group	0.67	0.50
F group	0.35	0.85
Experiment 1d (10-item lists, $N = 20$)		
R group	0.52	0.48
F group	0.30	0.54
Experiment 1e (10-item lists, $N = 36$)		
R group	0.48	0.35
F group	0.37	0.48

Results

As we planned to compare the differences between means as described above the analyses are organized around a series of *t* tests between pairs of means. However, for completeness we also report an ANOVA and the important groups × lists interaction but note that the *t* tests are conducted whether or not this interaction is significant following recommended practice for planned comparisons (Keppel, 1973/1991). Thus, the data were analyzed in a mixed-model ANOVA with groups as a between-subjects factor with two levels, R group and F group, and lists as a repeated measure also with two levels, list 1 and list 2; critical differences between means were evaluated with *t* tests. Table 1 shows the mean probabilities recall for groups and lists. The groups × lists interaction was significant, $F(1, 22) = 10.6$, $MS_e = .403$, $p < .01$, and, for the critical contrast of F group list 1 versus R group list 1 a significant difference was observed, $t = 4.9$, $p < .01$, indicating a robust DF effect. Also present was a difference between F-group list 1 performance and F-group list 2 performance in which that for F group list 1 was significantly poorer than that observed for list 2, $t = 2.6$, $p < .03$. Group differences in list 2 recall were not reliable, $t < 1$, although the means are in the predicted direction. Note that as this latter effect rarely reaches significance in any of the experiments described below it is not reported further. These findings then show a "standard" DF effect with F group list 1 recall reliably lower than R group list 1. In general there were a slightly lower levels of recall overall compared to other DF by lists studies (MacLeod, 1998) and, possibly, this may be related to the inclusion of the 5-min filled interval rather than more usual immediate recall. The filled interval, however, was used to minimize ceiling performance on list 2 which we had found in pilot studies preceding the present experiments.

In the DF procedure, and across the experiments reported here, comparisons are made between different groups of participants both within and between experiments. Of special interest is the comparison of F-group list 1 mean recall across experiments and, particularly, to the baseline mean in the present experiment. It is important then that this baseline F-group list 1 mean is in fact representative of performance in other DF experiments that are similar or which approximate in design to the present experiment. Table 1 shows the performance from five DF experiments conducted in our laboratory as part of other projects. Experiments 1a and 1b show data from healthy middle-age participants acting as controls to brain-damaged patients. In these two experiments list length was short and the midlist cue was within subjects with the F condition presented at the start of a session followed a series of neuropsychological tests, concluding with the R condition. In both experiments the condition × list interaction was significant, $F(1, 10) = 10.6$, $MS_e = .21$, $p < .01$ and $F(1, 13) = 17.9$, $MS_e = 1.01$, $p < .01$, respectively, as were the critical comparisons of F list 1 with R list 1, $t = 4.1$, $p < .01$ and $t = 3.3$, $p < .01$, indicating strong DF effects in both experiments, despite their departures from the standard design. In Experiment 1c different groups studied 20-item lists consisting of line drawings of common objects each presented for 2 s followed by immediate recall. A reliable interaction of groups × list was found, $F(1, 22) = 233.6$, $MS_e = 1.3$, $p < .01$ and the F list 1/R list 1 contrast was also significant, $t = 9.5$, $p < .01$. Experiments 1d and 1e approximate most closely to the "standard" design use here, the only variation being that participants in Experiment 1e were 14- to 15-year-old school children, whereas in 1d they were 19- to 21-year-old university students (as in Experiment 1c). In both Experiment 1d and 1e the group × list interaction was significant, $F(1, 18) = 8.3$, $MS_e = .19$, $p < .01$ and $F(1, 34) = 10.1$, $MS_e = .27$, $p < .01$, respectively, as were the critical comparisons of F list 1 with R list 1, $2.7 = 4.1$, $p < .02$ and $t = 2.2$, $p < .05$.

These findings demonstrate the robustness of the DF manipulation and also show how F list 1 mean recall values vary over experiments. If the data from each of the F list 1 conditions are entered in a one-way between-subjects ANOVA in which Experiments 1b through 1e are treated as groups, then there is no main effect of Experiment, $F = 1.1$. The overall mean probability of recall for F group list 1 is .36 with a *SD* of .16. The mean probability of F-group list 1 recall from Experiment 1 (.32) then falls well within half a *SD* of this overall mean from which it does not differ significantly, $t < 1$. Moreover, if the F-group list 1 comparison is made between Experiment 1 and Experiments 1d and 1e (which are highly similar in design to Experiment 1) then there are no reliable differences and the overall mean of Experiments 1d and 1e is .33, which compares favorably with the .32 mean from Experiment 1. We conclude that the value of F-group list 1 recall observed in Experiment 1 is a representative figure of DF in the "standard" procedure used in the present and later experiments and, conse-

TABLE 2

Free Recall of Lists 1 and 2: Mean Probabilities For Experiments 2 through 6

	List 1	List 2
Experiment 2: DF with vowel counting		
R group	0.73	0.32
F group	0.48	0.28
Experiment 3: DF with word length judgement		
R group	0.65	0.51
F group	0.50	0.41
Experiment 4: DF with concurrent memory load		
R group	0.62	0.48
F group	0.64	0.33
Experiment 5: DF with a semantic associate		
R group	0.67	0.49
F group	0.33	0.56
Experiment 6: DF with several semantic associates		
R group	0.67	0.6
F group	0.74	0.78

quently, constitutes an acceptable baseline against which to compare F-group list 1 performance in the experiments reported next.

EXPERIMENT 2

The second experiment followed the procedure used by Macrae et al. (1997), but with different stimuli. In this procedure participants in both groups kept a running total of the vowels in list 2 items and reported this prior to free recall of the lists. The only difference from the general method is that following the midlist cue participants were told that while learning second-list items they should count the vowels in each item and keep a running total that they had to report, in writing, on paper provided, as soon as the experimenter signaled that the list presentation was complete. After this the participants were given the arithmetic test and proceeded as in the general method.

Results

The analyses were the same as those in Experiment 1 and Table 2 shows the mean probability of recall. Note that all participants reported the correct total of vowels. The groups × lists interaction was significant, $F(1, 22) = 4.9$, $MS_e = .130$, $p < .04$, and the critical contrast of F group list 1 with R group list 1 was also reliable, $t = 4.4$, $p < .01$. Unsurprisingly, performance of the F group on list 1 was significantly higher than on list 2, $t = 3.7$, $p < .01$, and attenuated performance on list 2 for both groups had been expected given that the effect of the secondary task would have been to divide attention and so lower the degree and quality of learning of these items relative to list 1. Thus, the difference between the groups for list 1 shows the expected DF effect despite the secondary task on list 2. However, the F group's performance on list 1 is raised relative to the F group's performance on list 1 in Experiment 1. Comparisons between means found that the Experiment 1 F-group list 1 recall rate was significantly poorer than that observed in the corresponding condition in the present experiment, $t = 2.0$, $p < .05$. Note also that the two R groups in each experiment did not differ reliably in list 1 recall and that the level of recall of the TBF list in the present experiment at 48% is highly consistent with the level observed by Macrae et al. (1997) of 49%. In our view this finding demonstrates a *disruption* of inhibition that is brought about because the impoverished encoding of list 2 does not lead to a memory representation that triggers a powerful inhibitory response.

EXPERIMENT 3

In Experiment 2 we suggested that the effect of a secondary task on list 2 in the F group was to lead to a weaker inhibition of the TBF list 1 and we elaborate this explanation in some detail later. Here we want to consider how vowel counting has this attenuating effect upon inhibition and there are several components in the task that could, potentially, have weakened inhibition in the way observed. For instance, simply making some type of relatively undemanding judgment of the items in list 2 might be sufficient to weaken inhibition. On the other hand, holding in mind a concurrent memory load (the vowel total), which admittedly in the Experiment 2 only amounted to two digits, might be the critical factor. Yet another possibility is that only when two or more processing activities are combined together while learning list 2 is inhibition of a TBF list 1 disrupted. The present experiment investigates the first of these suggestions by having participants judge word length of list 2 items, a shallow processing task similar in this respect to vowel counting but entailing no concurrent memory load. Words were selected from the word pool to create two lists, each of which contained five words of six letters or less (short) and five words of seven letters or more (long). Within each list the short and long words were intermixed unsystematically with the constraint that runs of no more than two of one type of word were permitted. After the midlists cue participants were told that when a word went off-screen they should press one of the keys marked “S” for short or “L” for long. The “S” key was to be pressed if the word was six letters or less in length and the “L” was key was to be used if the word was seven letters or longer.

Results

Only a few errors were made in the word-length judgment task, amounting to less than 5% of total responses, and this indicates that participants in general performed this task to a high level. Table 2 shows the mean probabilities of recall in each of the four conditions and it can be seen that performance on list 2 was higher in this experiment than in Experiment 2, confirming that the word-length judgment task takes up less processing capacity than vowel counting. No reliable interaction was observed in the ANOVA of the free-recall data and the only term to reach significance was that of the main effect of lists, $F(1, 22) = 4.7$, $MS_e = .163$, $p < .05$, in which recall was higher in list 1 than 2. In the critical contrast, recall by the F group of list 1 was lower than that of the R group, but this difference was only marginally significant, $t = 2.13$, $p < .057$. This appears to be due to a fall in the R group’s recall of list 1 rather than a rise in recall by the F group of the inhibited list. Recall of list 1 by the F group in the present experiment and by the F group in Experiment 2 did not differ significantly, $t < 1$, but was reliably higher than recall of the F list in Experiment 1, $t = 2.22$, $p < .03$. Taken together this pattern of findings shows a disruption of inhibition of the F list very similar to that observed in Experiment 2 and this suggests that simply performing a second task while learning list 2 is sufficient to significantly weaken inhibition of the F list.

EXPERIMENT 4

The findings of Experiment 3 suggest that a concurrent memory load during second-list learning probably contributes little to disrupting inhibition. However, the concurrent memory load in Experiment 2 and in Macrae et al. (1997) was very light, amounting to no more than two digits. In the present experiment it was decided to use a digit list the length of which would fall at about span for our subject population (undergraduates) and we accordingly we used a six-digit list. After receiving the F-cue participants were, then, given a list of six randomly selected digits to keep in mind while learning list 2 and to report back as soon as the experimenter indicated that the list presentation was complete. The digit list was printed on a card placed facedown to the side of the computer and was turned faceup and studied for 20 s immediately prior to presentation of list 2. At the end of the 20-s period the card was placed facedown so that the digit list was no longer visible. At the end of list 2 presentation participants wrote down the digit list on the card and then turned immediately to the filler task.

Results

There were errors on the secondary task and several participants in both groups failed to recall the exact digit list. Nevertheless, even those participants who made errors reported six digits which were judged to correspond at least in part to the original list. Table 2 shows the mean probabilities of recall in each of the four conditions and it can be seen from this that list 1 items were recalled to a high level, whereas list 2 items were poorly recalled and this effect was reliable, $F(1, 22) = 22.7$, $MS_e = .585$, $p < .01$. This effect holds good for both the R and F groups and arises because the effects of

DF have been completely abolished. Indeed, recall of list 1 by the F group in the present experiment was significantly higher than recall of this list in any of the previous experiments. No other effects were significant, although the groups × lists interaction approached significance, $F(1, 22) = 3.6$, $MS_e = .09$, $p = .0725$, and this occurred because recall of list 2 in the F group was significantly lower than recall in all other conditions. This finding shows that holding a concurrent memory load in mind while trying to forget a previously acquired list may be especially difficult as acquisition of the second list was more severely hurt in the F than R group while at the same time inhibition of list 1 was abolished. It seems as though the processing resources required for learning and inhibition were all channeled into retaining the concurrent memory load. Indeed, all participants reported consciously rehearsing the digit list while studying list 2 and this in itself may have been sufficient to lower learning and abolish inhibition and we return to this point later.

Related to this was the outcome from an identical experiment which differed from the present experiment in that a nine-digit supraspan list was used rather than a six-digit at-span list (see Harries, 1999, for a full account of this experiment). Mean probabilities of recall in the nine-digit experiment for the R group were list 1 = .67 and list 2 = .44 and for the F group, list 1 = .53 and list 2 = .72. The group × lists interaction was significant, $F(1, 22) = 23.9$, $MS_e = .51$, $p < .01$, and F group list 1 differed reliably from F group list 2, $t = 2.53$, $p < .03$, but no other reliable contrasts for F group list 1 were observed. Also as recall of the F list in the nine-digit concurrent memory load experiment was poorer than recall of the corresponding list in the six-digit experiment, it seems that there may have been some inhibition of the TBF list in the nine-digit experiment, whereas there was no evidence of any inhibition whatsoever in the six-digit experiment (see Table 2). These findings, then, are highly inconsistent with those of the six-digit experiment but can be accommodated when it is appreciated that none of the participants in the F group were able to report back their nine-digit lists with anything but extremely low accuracy and, moreover, several of the participants in this group were unable to recall the list at all. Participants did not report conscious rehearsal of the digit list. The R group fared better, with all participants reporting nine digits, and seven of this group were judged to have moderately accurate recall; several of this latter group reported conscious digit-list rehearsal during list 2 acquisition. It seems then that a supraspan concurrent memory load cannot be held in mind while learning a second list and simultaneously attempting to inhibit a recently acquired list. Our F-group participants dealt with this overload by abandoning the nine-digit concurrent memory load task and, as a consequence, list 2 was acquired to a high level and some weak inhibition of list 1 then appeared.

EXPERIMENT 5

Experiment 4 demonstrated that inhibition can be abolished altogether with the appropriate list 2 secondary task. In the present experiment we investigate how inhibition may be overcome not without the use of a secondary task but rather by promoting integration of the lists. Golding, Long, and MacLeod (1994) found in an item-by-item DF experiment where some F-cued items were highly related to R-cued items that inhibition of F items associated with R items was disrupted and there was consistently higher recall of these associated items relative to the recall of F items unrelated to R items. Although, overall, recall of associated F items was still poorer than R item themselves. Thus, following Golding et al. (1994), it should be possible to considerably raise recall of the F list, while at the same time maintain a high level of recall of list 2 by using words on one list that are associated with words on the other list. The final two experiments explore this possibility. In the present experiment, which is identical to Experiment 1, a single word in the TBF list was highly associated with a single word in the TBR list. A set of pairs of associates were selected from word-association norms (Moss & Older, 1996) and for each pair only the two most closely associated pairs were used.

Results

Table 2 shows the mean probabilities of recall in each of the conditions and it can be seen that the pattern is that of a "standard" DF effect and virtually identical to that found in Experiment 1. There was a reliable groups $\times$ lists interaction, $F(1, 22) = 24.9$, $MS_e = .48$, $p < .01$, and F-group list 1 performance was reliably poorer than R group list 1, $t = 6.3$, $p < .01$, and F group list 2, $t = 4.2$, $p < .01$. Thus, when a single item on the TBF list is related to a single item on TBR list inhibition is unaffected.

EXPERIMENT 6

The results of Experiment 5 did not support our view that by increasing the similarity of the lists, albeit in a small way, inhibition would be weakened or even abolished. Given the results of Golding et al. (1994), this was surprising and we concluded that some small amount of similarity between the lists was not sufficient to affect inhibition. Consequently, it was decided to increase the number of associated pairs across lists to five word pairs (or half the list).

Results

Table 2 shows the mean recall probabilities and it is evident that recall in all conditions is high. Interestingly, the only manipulation to reach significance was a main effect of groups, $F(1, 22) = 9.1$, $MS_e = .2$, $p < .01$, and the F group benefited more from the presence of pairs of associated words than did the R group. A question of some interest that arises with these data is whether for the TBF list those items not related to list 2 TBR items also showed an improvement in recall relative to R group recall of the corresponding items from list 1; that is, whether the abolition of inhibition extended to unrelated items on the TBF list for the F group. Accordingly, a further ANOVA was conducted including all the factors analyzed previously but with the additional factor of related versus unrelated lists. A main effect of groups was found, $F(1, 22) = 10.8$, $MS_e = .14$, $p < .01$, and the mean for the R group was 62% compared to 78% for F group. A main effect of related versus unrelated item sets was also found, $F(1, 22) = 61.1$, $MS_e = .55$, $p < .01$, and mean recall of related items was 85% compared to 55% for unrelated items. No other effects were significant. Thus, F group recalled more overall than R group for both lists and, unsurprisingly, related pairs were recalled to a reliably higher level than unrelated pairs. In fact, recall of the F list was as high in this experiment as in Experiment 4 and reliably higher than in all other experiments, showing a complete abolition of inhibition.

TABLE 3

Recognition Collapsed over Lists 1 and 2: Mean Probabilities of Hits and False Alarms (FAs) for Experiments 1 through 6[a]

	R	K	F	Total
Experiment 1: "Standard" DF procedure				
Hits	0.57	0.21	0.11	0.89
FAs	0.04	0.02	0.14	0.20
Experiment 2: DF with vowel counting				
Hits	0.55	0.08	0.16	0.79
FAs	0.03	0.02	0.16	0.21
Experiment 3: DF with word length judgement				
Hits	0.64	0.11	0.13	0.88
FAs	0.02	0.04	0.13	0.19
Experiment 4: DF with concurrent memory load				
Hits	0.66	0.13	0.13	0.92
FAs	0.01	0.02	0.21	0.24
Experiment 5: DF with a semantic associate				
Hits	0.67	0.06	0.15	0.88
FAs	0.02	0.02	0.12	0.16
Experiment 6: DF with several semantic associates				
Hits	0.74	0.07	0.10	0.91
FAs	0.03	0.01	0.03	0.07

[a] R, Recollective Experience; K, Just Know; F, Feels Familiar.

RECOGNITION PERFORMANCE

The pattern of correct recognition performance and false alarms was the same in all six experiments and the mean probabilities are shown in Table 3. It can be seen from Table 3 that recollective experience is the dominant response category for hits in all experiments. The data for the hits were entered into a mixed-model ANOVA identical to that used for the the free-recall data, with the addition of memory awareness as a within-subjects factor with three levels, remember, know, and familiar, and this was done for each experiment. In each experiment there was a significant main effect of the variable memory awareness, $F(2, 44) = 16.6$, $MS_e = 2.8$, $p < .01$, $F(2, 44) = 34.5$, $MS_e = 3.03$, $p < .01$, $F(2, 44) = 100.1$, $MS_e = 4.3$, $p < .01$, $F(2, 44) = 37.4$, $MS_e = 4.5$, $p < .01$, $F(2, 44) = 79.1$, $MS_e = 5.3$, $p < .01$, and $F(2, 44) = 104.8$, $MS_e = 6.9$, $p < .01$, for Experiments 1 through 6 respectively. In several of the experiments there were higher order interactions but none of these showed any effects of inhibition on memory awareness (see Harries, 1999, for a full account). Thus, the prediction that correctly recognized TBF items might be differentially associated with familiarity responses was not supported. It is, however, important to note that the present study used the lists method of DF, whereas the Gardiner et al. (1994) study, from which our prediction was drawn, used the item method. This alone may be sufficient to account for discrepancies between the two studies. In the items DF procedure the TBF items are not considered to be inhibited, but instead they receive less elaborative rehearsal and it is this which makes them less memorable (MacLeod, 1998) and, as a consequence, more likely to be associated with feelings of familiarity when they are recognized (Gardiner et al., 1994). In the list method of DF the view is that the TBF list, originally acquired normally, is subject to inhibition following the F cue. The present findings demonstrate that recognition so completely overcomes this inhibition that the TBF items can be recollectively experienced during recognition, which lends strong support to the claim that the F list is, indeed, encoded normally and in a form appropriate to memorizing, i.e., in an elaborative way likely to facilitate recollective experience.[1]

ANALYSES OF OUTPUT ORDER

We have interpreted our findings in terms of inhibition caused by competition between lists 1 and 2. By this view the midlist F instruction defines list 1 as a potential competitor to list 2 and inhibition of the representation of list 1 is then triggered by presentation and during the learning of list 2. If, however, the second-list items are not encoded in such a way as to be strong competitors for list 1 items then only weak inhibition is triggered. Thus, when the encoding of list 2 is shallow or impoverished compared to the elaborative encoding of list 1 (Experiments 2 through 4) extensive inhibition is not triggered and performance on F list 1 rises—inhibition is disrupted. In contrast, an R midlist cue is an implicit instruction to integrate the two lists and, in this case, any automatic inhibition caused through list 2 learning is overcome by the integration of lists 1 and 2. There is, however, an alternative account of these findings which argues that the effects arise not from inhibitory processes operating during or close to encoding but, rather, from retrieval dynamics.[2] Several different versions of this account have been proposed (cf. Anderson & Spellman, 1994; and see Crowder, 1976; and Murdock, 1974; for reviews of earlier studies)

[1] Table 3 also includes the mean false-alarm rates which, in each experiment, also had reliable effects on the memory awareness variable. These are discussed in Harries (1999) and are shown here for completeness. We note that highest false-alarm rates occurs in the familiarity category and this is often the case in recognition memory experiments that take these awareness measures.

[2] Brainerd, Reyna, Harnishfeger, and Howe (1993) have shown in multitrial free-recall experiments that more difficult to retrieve items are recalled at earlier output positions. Whether this also applies to the single-trial recall used in the present experiments is not known; however, our suggestion is that it does not. The harder to retrieve items will always be from the F lists and if it is the case that in single-trial recall hard items are retrieved earlier then there should be a large number of participants, in the F groups, across all experiments, who systematically recall F items first. As no triage pattern (cf. Brainerd et al., 1993) was observed it seems reasonable to conclude that the findings from repeated recall studies do not extend to single-trial recall studies or at least the ones reported here. However, one possibility that might be noted is that earlier recall of harder items in multitrial free recall experiments might reflect the operation of retrieval inhibition of easier to recall items in later recall trials (cf. Roediger, 1974).

although the basic idea is simply that prioritizing recall from one list impairs recall (by whatever means) of the other list. So if, for example, the F group predominantly started their recall with list 2, this might cause an impairment of list 1 items which become progressively more inaccessible as more list 2 items are recalled (e.g., Roediger, 1973, 1974). Thus the inhibitory effects in the DF procedure could arise because of order of output at recall and not because of earlier inhibition of list 1. There are various arguments and findings, which suggest that this alternative is incorrect; but before turning to those we report a final study that aimed to rule out the output account.

EXPERIMENT 7

The output hypothesis as stated above can be tested by requiring recall of list 1 first. If, under these circumstances the DF effects in the standard procedure and in the list 2 with secondary tasks procedure are still observed, then it could not be maintained that the effects arise from output interference. For instance, if list 1 in the F group shows an impaired level of performance similar to that observed in Experiment 1 and the other standard DF experiments reported earlier, then the impaired performance could not be attributed to the interfering effects of recalling list 2 first. Similarly, if the level of F group list 1 recall following vowel counting on list 2 but with recall of list 1 first is comparable to that observed in Experiment 2, then, again, the pattern of disrupted inhibition could not be attributed to output effects. In fact there are good reasons to suppose that this latter pattern will in fact be the outcome as list 1 was spontaneously recalled first by most participants in Experiment 2 (see analyses of output order below). Accordingly, then, we repeated Experiments 1 and 2 with the following change: at test participants were instructed to recall items from list 1 first and only when they could recall no more were they then to recall items from list 2.

Results

Table 4 shows the mean probabilities of recall in each of the four conditions in a standard DF procedure (top panel) and with vowel counting on list 2 (middle panel). Note that there were 15 participants per group. Turning first to the standard procedure, the groups × lists interaction was significant, $F(1, 28) = 32.4$, $MS_e = 7.4$, $p < .01$, and for the critical contrast of F group list 1 versus R group list 1 a significant difference was observed, $t = 7.4$, $p < .01$, indicating a DF effect. The similarity of these findings to those of Experiment 1 (see Table 1) suggests that output interference, as conceived of here, is not a factor which strongly influences the pattern of recall.[3] The major difference between the two sets of data is an overall increase in the level of recall in all four

TABLE 4

Free Recall of Lists 1 and 2: Mean Probabilities for Experiment 7 and a Subset of Experiments 1 and 5

	List 1	List 2
Experiment 7: Standard DF procedure with recall of list 1 first		
R group	0.81	0.58
F group	0.40	0.62
DF with recall of list 1 first and vowel counting on list 2		
R group	0.72	0.35
F group	0.49	0.32
DF for participants recalling list 1 first in Experiments 1 and 5		
R group ($n = 13$)	0.67	0.47
F group ($n = 7$)	0.40	0.49

[3] In a third replication of Experiment 1 with recall of list 1 first we found, however, a level of .48 for F list 1 recall. This comparatively high level of recall was significantly different from Experiment 1 F list 1 recall, $t = 2.08$, $p < .05$, but did not differ reliably from the level of recall in this condition in the present experiment, $t = 1.3$. A closer examination of the data from this third replication identified three outliers in the F group, with near-ceiling scores of .8 on list 1 and .9 on list 2; by comparison the next highest overall score by any participant in this group was .6. If the scores of these outliers are removed and the data reanalyzed, then all the effects remain significant, the only change being that F list 1 mean performance falls to .4, which is not significantly different from the corresponding performance in Experiment 1 and is identical to the level reported in Table 4. It seems to us that these data are interesting because they (tentatively) suggest that there may be strong individual differences in overcoming inhibition—an issue which has not yet been examined with the DF procedure.

conditions in the present experiment compared to Experiment 1 and this may arise because the increased difficulty of the recall task induces more effort. Note that the F-group list 1 performance is higher here (Table 4) but does not differ reliably from the same level of performance in Experiment 1, $t = 1.1$. The replication of Experiment 2 (middle panel of Table 4) also produced a reliable groups × lists interaction, $F(1, 28) = 7.4$, $MS_e = 1.6$, $p < .02$, with a significant difference between groups on list 1, $t = 4.4$, $p < .01$. It can be seen from Table 4 that recall of list 2 was hurt for both groups and this simply shows that vowel counting impairs retention of the list. Directing recall to list 1 first had no noticeable effect in either group and the means shown in Table 4 are virtually identical to those previously observed in Experiment 2 (see Table 2). Taken together the findings of these two replications would seem to rule an output account of the findings.

A further way in which to test the output hypothesis is to examine DF effects in those participants who spontaneously recalled list 1 items first. If reliable DF effects are present here, then this too would also strongly argue against the output hypothesis. Moreover, data from participants who spontaneously recall list 1 items first have the advantage of being a pure index of the effects of output order in that attention has not been explicitly directed to this list. The data for Experiments 1 and 5 were reexamined[4] and those cases where the first item to be output was from list 1 were selected for further analysis. Note that recall of the first item from list 1 was used as a criteria because by recalling a list 1 item first the probability of recalling other list 1 items is increased (assuming no inhibition). Thus, if output interference is the operative factor in the DF effect then when list 1 items are recalled first this should interfere with recall of list 2 items and, consequently, no DF effect will be observed. Six participants from Experiment 1, three in the F group and three in the R group, recalled an item from list 1 first and in Experiment 5 14 cases were identified, 4 in the F group and 10 in the R group. These data were pooled to form a group of 7 cases of F group recalling list 1 first and 13 cases of R group and the mean probabilities are shown in Table 4. It can be seen from Table 4 that F group list 1 recall is poorest overall and the groups × lists interaction is significant, $F(1, 18) = 6.7$, $MS_e = .19$, $p < .02$, as is the critical contrast of F group list 1 with R group list 1, $t = 2.94$, $p < .01$, indicating a robust DF effect. The difference between the F group list 1 mean of .40 observed here and the corresponding condition in Experiment 1 (with the six participants removed) was not significant; moreover this value is identical to that obtained earlier (top panel of Table 4). It seems, then, that output interference contributes little to the effects reported above and no matter which list is recalled first powerful DF effects are observed. When attention is explicitly directed to an inhibited list then inhibition for one or two items may be overcome but despite this, performance on the inhibited list remains significantly depressed compared to the control condition.

Experiment 7 and the additional analyses of Experiments 1 and 5 effectively rule out an account of our findings in terms of output interference; nevertheless it will be useful to have an overall index of output order that allows comparisons between conditions and experiments. In order to achieve this the ranked output order of items from each list for each participant was tabulated. A scaled rank for each list was computed for each participant in the following way: the mean rank for all list 1 items (MR_1) of an individual participant was calculated. Scaled rank was then computed by subtracting $(N1 + 1)/2$ from MR_1 (where $N1$ = number of list 1 items recalled) and the product divided by $N2$ (number of list 2 items recalled). In calculating the scaled ranks for list 2 MR_2 replaced MR_1 and $N2$ replaced $N1$. The advantage of this transformation is that it takes into account both the number of items recalled from a list and their various ranks and expresses the scaled score as a number between 0 and 1. A scaled score of .5 indicates that the items were interspersed equally in recall, i.e., output alternated from list 1 to list 2, a score in the direction of 0

[4] Data from these two experiments were used because they produced a standard DF effect.

means that those items tended to be recalled first, and a score in the direction of 1 that those items tended to be recalled last. For each of the experiments, one through six, the scaled ranks (SRs) were entered into an ANOVA identical to that used with the free recall data.

In Experiment 1 a reliable groups × list interaction was found, $F(1, 22) = 7.6$, $MS_e = 1.74$, $p < .01$. The R group showed a bias to output list 1 items first, SR = .3, and to list 2 items second, SR = .7. In contrast, the F group bias was to list 1 items second, SR = .7, and to list 2 items first, SR = .3. As Experiment 7 demonstrates that in the standard DF procedure F group list 1 performance remains depressed even when all list 1 items are recalled it follows that output order is unrelated in level of recall, at least in the present experiments. In fact this striking interaction in output order is exactly the pattern that would be expected if the most highly active or accessible items are recalled first. Thus, in the F group list 1 is inhibited and, therefore, has low accessibility, whereas list 2, which is not inhibited, has high accessibility, ensuring early recall of items from this list. The position with the R group is more complex and an account proposing that, for example, retroactive interference from list 1 to list 2 leads to the former list obscuring or somehow overshadowing the latter and, hence, biasing recall of list 1 first, although plausible, is not especially compelling. Instead, it might just as well be proposed that all else being equal, participants when recalling two recently learned lists have a general bias to start their recall with the items learned first. The problem with all these explanations is that in Experiment 5, where there was also a standard DF effect, no reliable groups × lists interaction was found and output SRs here were close to .5, indicating interspersing of items from the two lists in order of recall. In contrast, SRs from Experiment 2, L1 = .3 and L2 = .7, $F(1, 19) = 12.9$, $MS_e = 1.9$, $p < .01$; Experiment 3, L1 = .3 and L2 = .7, $F(1, 22) = 13.5$, $MS_e = 1.9$, $p < .01$; and Experiment 4, L1 = .4 and L2 = .7, $F(1, 22) = 5.5$, $MS_e = 1.1$, $p < .03$, all showed reliable list effects with biases to recalling list 1 items first. This, however, is unsurprising, as in each case list 2 was acquired under conditions of divided attention and consequently would have received only a shallow encoding. Finally, in Experiment 6 only lists was reliable, L1 = .4 and L2 = .6, $F(1, 22) = 5.3$, $MS_e = .5$, $p < .01$; and despite this bias to recall list 1 items first the SR values indicate extensive interspersing of items during recall.

Finally, one potential problem with the SR method of calculating biases in output order is that it does not take into account amount recalled. Thus, a participant who recalled two list 1 items followed by two list 2 items would have SRs of 0 and 1 respectively, as would a person who recalled four list 1 items first followed by eight list 2 items. This is not a problem if the interest is simply in biases in output *regardless of the amount recalled.* However, it might be argued that a participant who recalled, say, two list 1 items first followed by eight list 2 items was, in fact, subject to output interference of list 1 by the extended recall of list 2 items.[5] Such output interference should be present in the R and as well as F groups and the R groups are good candidates to show this as their recall is not subject to an earlier F instruction. Accordingly the R group data from Experiments 1 and 5 were examined and six participants with SRs of 0 (all list 1 first) and 1 (all list 2 second) were identified (there were also two cases of participants with SRs of 1 and 0). If it is the case that later and extensive recall of list 2 interferes with the further recall of list 1 then there should be greater recall of list 2 items than list 1 items in this subgroup of participants. It was found that mean list 1 recall was .76 compared to list 2 recall, which had a mean of .63. These, however, did not differ reliably, $t < 1.6$, and this suggests that "overshadowing" output inference was not a systematic factor in the present DF experiments. In fact, looking over all participants' patterns of recall in Experiments 1 and 5 there were no cases which could, on the basis of output order, be unambiguously classified as a product of output interference. Patterns of recall which did not yield SRs of 0/1 or 1/0 were of

[5] We thank Mike Anderson for bringing this alternative account to our attention.

mixed form such as 11222111121211, 112221222111, and so on, and this is reflected in the mean SRs, which, with values ranging between .3 and .7, indicate mixed patterns of recall in which lists were sampled several times at different points during recall.

In summary, the present section shows that output interference played little if any role in the patterns of recall reported in Experiment 1 to 6 and this, perhaps, is not so surprising. University students have few problems in recalling all of two lists of 10 words each. In our pilot studies we found that we had to introduce a demanding filled interval in order to prevent the R groups performing at ceiling and even with this R group list 1 performance averaged .67 over the first six experiments. It seems to us that failure to recall from F group list 1 was largely a product of inhibition with, possibly, output interference affecting the accessibility of the occasional item for some participants. Finally, this conclusion is consistent with other findings, for example, by Geiselman, Bjork, and Fishman (1983), who found that directing recall to commence with one list or the other does not remove the DF effect, again showing that output order is not critical. More recently it has been found that when participants recall only *one* list, i.e., list 1 or list 2 in both F and R groups, the DF effect remains robust (R. Bjork, personal communication). Taken together the experiments and additional analyses reported in this section rule out an account of the overall pattern of findings in terms output order. This does not mean that output order effects did not occur in the experiments, nor does it mean that retrieval dynamics do not effect level of performance in other recall tasks; instead, our claim is that output order is not a factor that systematically influences level of performance in the DF by lists procedure, and derivatives of it, used in the present sequence of experiments.

GENERAL DISCUSSION

In summary, we have found that inhibition of a recently acquired list of words can be disrupted if a secondary task is performed while learning a second TBR list of words (Experiments 2 and 3) and can be completely abolished if a within-span concurrent memory load accompanies second-list learning (Experiment 4). Abolition of inhibition also occurs when the two lists are strongly associated (Experiment 6), but not when they are only weakly associated (Experiment 5). None of these manipulations, on the other hand, whether inhibition inducing (Experiments 1 and 5) or inhibition disrupting (all other experiments), systematically influenced recognition performance, which generally was at a high level. Correct recognition was dominated by recollective experience, which contrasts with recognition in item-by-item DF experiments where an increase in feelings of familiarity characterizes correct recognition (Gardiner et al., 1994). We believe these differences reflect lack of elaborative rehearsal of F items in item-by-item DF experiments which gives rise to memory representations that promote feelings of familiarity rather than recollective experience during correct recognition. By contrast, in DF list experiments the F list is learned normally, i.e., as a TBR list, and because of this is, presumably, elaborately encoded. Consequently, when an item on a recognition test overcomes the inhibition the recognized word can be recollectively experienced. Finally, an extensive investigation (Experiment 7) and examination of the patterns of output order indicated that this was highly unlikely to have led to the patterns of recall observed in the present DF experiments. We turn next to an account of the overall pattern of findings, which in our view can best be explained in terms of how effective list 2 is in triggering inhibitory processes.

Inhibition, Integration, and Competition

A fundamental assumption that underlies our reasoning is that encoding processes aim to minimize the load imposed by the formation of new memories, while simultaneously keeping available, for a limited time, a large amount of recently acquired information. We believe that a key feature of experience to which inhibitory and integration encoding process respond is the *degree of potential for interference in later recall* of successively encoded episodes. When it is judged (and we assume such judgments are nonconscious) that temporally contiguous episodes are highly likely to interfere with each other in later recall, then each newly encoded episode is subject to inhibition, i.e., the encoding system effectively self-generates an F cue. The inhibition is triggered by the encoding of a subsequent competitor episode and it reduces the accessibility rather than the availability of the competitor episode already represented in long-term memory. Another way in which encoding processes deal with the uncertainty of what and how much should be retained is by integration. In this case, competing, or potentially competing, memory episodes are integrated with each other to form representations that do not trigger inhibitory encoding processes: integrated episodes do not compete because they are part of the same knowledge structure (cf. Radavansky, 1999). Although encoding processes may be designed so as to generally promote integration where this is possible, it may be that they too require some type of initiating cue, e.g., an R cue. Of critical importance to initiation of inhibition and integration are interepisode remindings that occur during encoding of a new episode. If the episode being encoded is highly similar to a recently encoded episode then *during the process of encoding* it will cue the recall of items from that episode. If, for whatever reason, the already encoded episode is to be inhibited then these interepisode remindings will act to intensify the inhibition. If, on the other hand, and again for whatever reason, the already-encoded episode is to have it's accessibility maintained or even increased, then the interepisode remindings will increase integration of the already-encoded episode with the episode currently being encoded (Experiment 6).

In order to apply this reasoning to the complex pattern of our findings we draw further upon a proposal by Anderson and Spellman (1994; see too Bjork et al., 1998; and Anderson & Bjork, 1994); namely that the process of retrieval inhibition is similar in type to the inhibition that occurs in selective attention. In selective attention it has been shown that directing attention to a specific target causes the simultaneous inhibition of competitor targets (e.g., Tipper, 1985; Tipper et al., 1991; see Pashler, 1998, for a review). Extrapolating from this, a simple and powerful explanation of the DF effect in Experiments 1 and 5 is that when attention is directed to the encoding of the second list, memory for the first (competing) list becomes a target for inhibition. If the similarity between the two lists is such that they would later compete in recall, then inhibition will be triggered and, moreover, will increase as items from the first list are occasionally accessed during list 2 learning. Thus, a critical factor in the DF effect is the potential of lists 1 and 2 to compete in later recall. This may also explain why an F cue on its own, i.e., without further list learning, has little impact on memory despite a deliberate and effortful attempt to forget. Indeed, this type of willed forgetting may have quite the opposite, and ironic, effect of increasing remembering of TBF materials (Wegner, 1994). Inhibition does not occur in the R groups because the first list is associated with an R cue and this triggers integration. Here, when list 1 items are occasionally accessed during second-list learning they will be associated with the list 2 item currently being processed and integration will be facilitated. Possibly, this is why we often observed a pattern of list switching in the output order data in the R groups.

Why then should a concurrent task (Experiment 3) on list 2 disrupt inhibition and a concurrent memory load (Experiment 4) on list 2 abolish inhibition? We propose that the effect of a secondary task during second-list learning is to reduce the extent to which the second list is encoded as a competitor to the first list. This occurs because the second list cannot be encoded in a rich enough way to act as powerful competitor to the representation in memory of list 1 (which has been elaboratively encoded). The secondary task on list 2 may also impair or prevent recall of list 1 items during list 2 learning, i.e., by drawing upon any unused processing capacity, and this may additionally act to lower list similarity and, hence, weaken inhibition. A problem for this account arises when the findings of Experiment 2 and 4 are compared. Recall of list 2 by the F group in both experiments is at about the .3 level (see Table 2), suggesting quite severe impairment in the acquisition of this list. According to our reasoning the ability of the representations of lists 2 and 1 to compete in recall should therefore be comparatively low and as a consequence recall of list 1 should be raised above the level of that observed in Experiment 1 *by about the same amount in both experiments*. But this was not what we observed and, instead, F-group recall of list 1 in Experiment 2 was 16% *less* than in Experiment 4—a significant difference. In spite of that result, we believe that our account in terms of competitors triggering inhibition remains viable. In our view, the difference between these two experiments depends on the way in which the different list 2 manipulations in each experiment impact upon list 2 learning and, consequently, determine list competitiveness. Vowel counting (Experiment 2) and a within-span concurrent memory load (Experiment 4) may impair second-list learning to about the same extent while impacting upon inhibition to different degrees. The particular relevant difference here is that the two-digit concurrent memory load in vowel counting is less demanding than that in the six-digit concurrent memory load used in Experiment 4. Our view is that in Experiment 4 list 1 items were rarely, if at all, accessed during list 2 learning and it is this that leads to the abolition of inhibition in this experiment. In contrast, the less demanding task of vowel counting may have allowed some access of list 1 items during list 2 learning and so weakened rather than prevented inhibition. Thus, inhibition can be disrupted in at least two way: either by a reduced competition or by prevention of access to list 1 items during list 2 learning. When the two are combined, as in holding a concurrent memory load in mind during second-list learning, there is a cumulative effect and inhibition is abolished (Experiment 4).

Finally, consider the results of Experiments 5 and 6. Introduction of a weak association between the lists in the F group does not affect DF and the usual inhibitory pattern is observed (Experiment 5). In contrast, introduction of a stronger association between list 1 and 2 items, when half of each of the two lists are associates of each other, completely removes inhibition. The pattern of mean free recall in Experiment 6 (see Table 2) showed that when there was a strong association between the lists not only was inhibition abolished but the F group benefited significantly more than the R group from the presence of interlist associates. We suggest that two factors played a role in producing this pattern of results. First, inhibition was abolished because the strong associations between the lists powerfully cued integration processes by promoting extensive recall of list 1 items while learning list 2. The net effect of this high degree of interlist cueing was that the two lists were treated (by encoding processes) as one and so, as they were no competitor lists, inhibitory processes were not triggered. Second, the increased recall in the F group, compared to the R group, may have been a type of "rebound" effect (Wegner, 1994). This latter effect may have arisen as a "correction" process to a misplaced attempt at inhibition. Assume that there is an intention to forget, induced by the F cue, and that presentation of the second list leads to identification of the memory of list 1 as a to-be-inhibited competitor. Inhibition is then highly prime and a target has been identified, but then the content of the second list turns out to be unexpectedly related to the memory of List 1. This association redefines list 1 not as a competitor, but as part of the same list and—as a consequence—additional activation, sufficient to overcome any initial inhibition, is directed to encoding the two lists and their interlist associations. Possibly, this additional activation is responsible for the enhanced recall of F group compared to R group in Experiment 6.

Our view, then, arising from the data reported above and derived from the general view of Bjork et al. (1998), is that inhibition is primed when one set of materials, already represented in memory, is defined as a potential competitor to another about-to-be-encoded set of materials. Inhibitory processes are then triggered or initiated by the encoding of the new TBR list. The strength of the inhibition is set by the (potential) degree of competition between the already-encoded TBF list and the TBR list currently being acquired. The more the potential of the TBF materials to compete later in recall with the TBR materials the greater the degree of inhibition. However, there is a trade-off between the degree of potential competition and list discriminability. Lists that are highly similar, because they contain many associates as in Experiment 6, may not be encoded separately, in which case inhibition would not be triggered because there would be no competitor memory representation. One strong prediction of the the view we have developed here is that it should be possible to depress the level of inhibition with sufficiently dissimilar lists and work in progress in our laboratory suggest that this is indeed the case (Racsma'ny & Conway, 1999). In a "standard" DF experiment the two lists comprised either all members of the same category, e.g., all vehicles, or were a mixed list of unrelated items. When the categorized list was used as list 2, F group mean recall of list 1 was .48 and although significantly less than recall of list 2, which was .64, it was significantly higher than standard F-group list 1 recall of .32. These more recent data and the theoretical account developed in this section suggest that complex DF results can be usefully conceptualized in a framework that focuses on how inhibitory and integration processes operate during encoding.

Shaping Memories

A further aim of the present experiments was to initiate development of an account of the neglected topic of the encoding of autobiographical memories (AMs). Our theoretical starting point is that in the short term, over periods of minutes and hours, very detailed memories are retained. Over longer retention intervals of days, weeks, months, and years, memories lose much of their detail, although they may nonetheless retain some sensory-perceptual details (Conway, 1996, 1999). Thus, participants in the present experiments might, after an extended retention interval of, say, several years, recall that they took part in some research when they were at university. Perhaps they could recall some details of the laboratory, a specific experimenter, and they may even remember that they had to learn lists of words. It is unlikely at this point that they will recall, as they once could, any of the words from the lists or that there had been two lists, that they received a midlist instruction, that there was a filled interval, a recognition test, and so forth. In general, the accessibility of details of AMs changes over time as memories are "shaped" by consolidation processes, such as rehearsal, into particular patterns. Note that we are not suggesting that this is the only way in which AMs take on their form and the predictability of experience (Schank & Abelson, 1977), as well as it's unpredictability (Kahneman & Miller, 1986), extent of prior knowledge, self-relevance, emotions, as well as a host of other factors (cf. Mullin, Herrmann, & Searleman, 1993), operating in concert will all influence what is retained in both the short and long term (Conway et al., 1996, 1997). Instead our point is that inhibition and integration processes operating at encoding may be one of the sets of factors that prepare a memory to have, as it is encoded or close to encoding, a particular pattern of accessibility and this pattern may, by influencing what aspects of an AM are subsequently rehearsed, may make an important contribution to determining the long-term "shape" of a memory. We believe that the DF procedure represents a good laboratory analog of the shaping of memories that takes place in everyday cognition and, therefore, constitutes an effective way in which to examine initial retention of AM details.

In the laboratory the "units" of experience to be sampled by a memory test are defined by the experimenter. In an analogous way the structure of experience, perhaps determined by the beginnings and endings of distinct actions, thoughts, and/or feelings (Conway, 1992, 1999; Newtson, 1976), may provide the junctures that define units of to-be-encoded experience in everyday cognition—much as the between-lists cue defines the lists in a DF experiment. So for example, switching from thinking to writing, from writing to making a cup of coffee, and so on, may all provide junctures at which units of experience become encoded into memory. These "units," which are summaries of epochs of ongoing dynamic cognition, contain extensive sensory-perceptual information—they are "near-experience"—and Conway (1996, 1999) refers to them collectively as Event-Specific Knowledge (ESK). The lists in a DF experiment can be viewed as analogous to ESK units (indeed, depending on how they are encoded, the lists may in fact be individual ESK units) and can be subject to retrieval inhibition in a similar way. Thus, as each unit is represented it competes, more or less strongly, with recently represented units. Patterns of inhibition then build-up that reflect the degree of competition. (Note that we assume that no conscious intention to forget is needed to produce this spontaneous retrieval inhibition, which is simply an automatic encoding process. A conscious intention to forget is only required in the artificial DF procedure.)

The outcome of the build-up of inhibition is a lowering of the accessibility of ESK units of a similar type up to the most recently encoded unit. By this view, inhibition is the default state of most units of ESK of a similar type with the exception of the most recent unit to be encoded. As a consequence, ESK units of a similar type remain available but have attenuated access. Unless these units are accessed, rehearsed, and integrated with autobiographical memory knowledge structures they may become permanently unavailable (forgotten). Thus, inhibition of similar ESK units shapes a memory so that only certain units are readily accessible in free recall. In addition to such effects, of course, are the current goals of the rememberer, which may override automatic inhibition and make accessible specific ESK units or groups of units that would otherwise have been inhibited (cf. Conway & Pleydell-Pearce, 2000). Personal goals may act as unifying themes that integrate ESK units and, in so doing, make them immune to inhibition, or at least the type of inhibition that underlies the DF effect. Thus, retrieval inhibition of competitor units takes place only for those units not subject to other intervention.

In our view the purpose of retrieval inhibition in encoding is to promote forgetting and, in particular, the forgetting of redundant, similar ESK units. This forgetting is achieved by making such units less accessible to recall and so less available for (spontaneous) rehearsal and subsequent integration with other autobiographical knowledge. Lack of rehearsal and integration will render a unit progressively more and more unavailable until at some point during the retrieval interval the unit will become permanently unavailable and, in effect, forgotten. These inhibited units of ESK can, however, be accessed by a sufficiently specific cue, as in recognition, and retrieval inhibition only influences what is accessed in recall. This latter aspect of retrieval inhibition allows the reevaluation of an event, or aspect of an event, when a subsequent cue overcomes the inhibition of a detail of an earlier memory. (How long the accessing of an inhibited ESK unit with a specific cue remains a possibility is unknown, but see Conway, 1997, for a review of data that suggesting it may remain possible for very long periods of time.) In conclusion, by shaping what can be easily recalled, retrieval inhibition indirectly influences which specific details might potentially be retained in accessible form over lengthy retention intervals while at the same time placing other (TBF) units in a state of attenuated accessibility that promotes their forgetting.

Finally, consider the outcome of encoding in conditions where retrieval inhibition is either attenuated or prevented altogether. One possibility is that the types of memories that result from Post Traumatic Stress Disorder (PTSD) are, at least in part, a product of a widespread failure of retrieval inhibition at encoding.[6] This failure minimizes the organizing influence of retrieval inhibition at encoding and allows the retention of apparently "irrelevant" minutiae (Brewin, 1998; Ehlers & Steil, 1995; van der

[6] Failure of inhibition at encoding may occur when markedly different ESK units are encoded consecutively, as discussed earlier. In the experience of a trauma it may be that the initiating event leads to an orienting response and as the event unfolds over time powerful emotions are generated. If this leads to the encoding of two separate units of ESK, one representing the period of the orienting response and the other period when emotions were experienced, then the latter representation may not compete with former. This would leave the AM for the orienting response highly accessible. Of course, other factors must also be operative but the failure of inhibition may help promote formation of an intrusive memory.

Kolk & Fisler, 1995). The sensory-perceptual qualities of PTSD memories are intense and they have a "reliving" property that makes their recall disturbing and destabilizing; this type of sensory "overload" is what might be expected when retrieval inhibition fails. Similarly the intrusiveness and repetitive recall of trauma memories may arise because they contain an excess of highly accessible redundant details (ESK units), making them particularly vulnerable to activation by specific cues (cf. Conway, 1997; Conway & Pleydell-Pearce, 2000). Retrieval inhibition may then be a basic organizing process in the construction and consolidation of durable and structured autobiographical memories and when it fails disorganized memories dominated by sensory perceptual details are the result.

REFERENCES

Anderson, M. C., & Bjork, R. A. (1994). Mechanisms of inhibition in long-term memory: A new taxonomy. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 265–325). San Diego, CA: Academic Press.

Anderson, M. C., & McCulloch, K. C. (1999). Integration as a general boundary condition on retrieval-induced forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* in press.

Anderson, M. C., & Spellman, B. A. (1995). On the status of inhibitory mechanisms in cognition: Memory retrieval as a model case. *Psychological Review,* **102,** 68–100.

Bjork, E. L., & Bjork, R. A. (1996). Continuing influences of to-be-forgotten information. *Consciousness & Cognition,* **5,** 176–196.

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

Bjork, R. A., Bjork, E. L., & Anderson, M. C. (1998). Varieties of goal-directed forgetting. In J. M. Golding & C. M. MacLeod (Eds.), *Intentional forgetting: Interdisciplinary approaches* (pp. 103–137). Mahwah, NJ: Erlbaum.

Brainerd, C. J., Reyna, V. F., Harnishfeger, K. K., & Howe, M. L. (1993). Is retrievability grouping good for recall? *Journal of Experimental Psychology: General,* **122,** 249–268.

Brewin, C. R. (1998). Intrusive autobiographical memories in depression and post-traumatic stress disorder. *Applied Cognitive Psychology,* **12,** 359–370.

Conway, M. A. (1992). A structural model of autobiographical memory. In M. A. Conway, D. C. Rubin, H. Spinnler, & W. A. Wagenaar (Eds.), *Theoretical perspectives on autobiographical memory* (pp. 167–194). Dordrecht, The Netherlands: Kluwer.

Conway, M. A. (1996). Autobiographical memories and autobiographical knowledge. In D. C. Rubin (Ed.), *Remembering our past: Studies in autobiographical memory* (pp. 67–93). Cambridge: Cambridge Univ. Press.

Conway, M. A. (1997). Past and present: Recovered memories and false memories. In M. A. Conway (Ed.), *False and recovered memories* (pp. 150–191). Oxford: Oxford Univ. Press.

Conway, M. A. (1999). Phenomenological records and the self-memory system. In T. McCormack & C. Hoerl (Eds.), *Time and memory: Issues in philosophy and psychology.* Oxford: Oxford Univ. Press.

Conway, M. A., Collins, A. F., Gathercole, S. E., & Anderson, S. J. (1996). Recollections of true and false autobiographical memories. *Journal of Experimental Psychology: General,* **125,** 69–95.

Conway, M. A., Gardiner, J. M., Perfect, T. J., Anderson, S. J., & Cohen, G. M. (1997). Changes in memory awareness during learning: The acquisition of knowledge by psychology undergraduates. *Journal of Experimental Psychology: General,* **126,** 1–21.

Conway, M. A., & Pleydell-Pearce, C. W. (2000). The construction of autobiographical memories in the Self Memory System. *Psychological Review,* in press.

Crowder, R. G. (1976). *Principles of learning and memory.* Hillsdale, NJ: Erlbaum.

Ehlers, A., & Steil, R. (1995). Maintenance of intrusive memories in post traumatic stress disorder: A cognitive approach. *Behavioural and Cognitive Psychotherapy,* **23,** 217–249.

Gardiner, J. M., & Parkin, A. J. (1990). Attention and recollective experience in recognition memory.

Gardiner, J. M., Gawlik, B., & Richardson-Klavehn, A. (1994). Maintenance rehearsal affects knowing not remembering: Elaborative rehearsal effects remembering not knowing. *Psychonomic Bulletin & Review,* **1,** 107–110.

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.

Harries, K. (1999). *Disruption and dissolution of directed forgetting.* Doctoral dissertation, Department of Experimental Psychology, Centre for Learning and Memory, University of Bristol, England.

Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review,* **93,** 136–153.

Keppel, G. (1991). *Design and analysis. A researcher's handbook.* Englewood Cliffs, NJ: Prentice Hall.

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

Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Ford,

R. L. (1997). On the regulation of recollection. The intentional forgetting of stereotypical memories. *Journal of Personality and Social Psychology,* **72,** 709–719.

Moss, H., & Older, L. (1996). *Birkbeck word association norms.* Hove, Sussex: Psychology Press.

Mullin, P. A., Herrmann, D. J., & Searleman, A. (1993). Forgotten variables in memory theory and research. *Memory,* **1,** 43–64.

Murdock, B. B. (1974). *Human memory: Theory and data.* Potomac, MD: Wiley.

Newtson, D. (1976). Foundations of attribution: The perception of ongoing behaviour. In J. H. Harvey, J. W. Ickes, & R. F. Kidd (Eds.), *New directions in attribution research* (Vol. 1, pp. 41–67). Hillsdale, NJ: Erlbaum.

Pashler, H. E. (1998). *The psychology of attention.* Cambridge, MA: MIT Press.

Racsma'ny, M., & Conway, M. A. (1999). Reducing inhibition in directed forgetting. In preparation.

Radvansky, G. A. (1999). Memory retrieval and suppression: The inhibition of situation models. Under review.

Roediger, H. L., III (1973). Inhibition in recall from cueing with recall targets. *Journal of Verbal Learning and Verbal Behaviour,* **12,** 644–657.

Roediger, H. L., III (1974). Inhibiting effects of recall. *Memory and Cognition,* **2,** 261–269.

Schacter, D. L., Norman, K. A., & Koustaal, W. (1998). The cognitive neuroscience of constructive memory. *Annual Review of Psychology,* in press.

Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals, and understanding.* Hillsdale, NJ: Erlbaum.

Tipper, S. P. (1985). The negative priming effect: Inhibitory effects of ignored primes. *Quarterly Journal of Experimental Psychology,* **37A,** 571–590.

Tipper, S. P., & Driver, J. (1988). Negative priming between response modalities: Evidence for the central locus of inhibition. *Perception and Psychophysics,* **43,** 45–52.

Tipper, S. P., Weaver, B., Camerson, S., Brehaut, J., & Bastedo, J. (1991). Inhibitory mechanisms of attention in identification and localization tasks—Time course and disruption. *Journal of Experimental Psychology: Learning, Memory and Cognition,* **17,** 681–692.

Tulving, E. (1983). *Elements of episodic memory.* Oxford: Clarendon.

Tulving, E. (1985). Memory and consciousness. *Canadian Psychologist,* **26,** 1–12.

Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behaviour,* **5,** 381–391.

van der Kolk, B. A., & Fisler, R. (1995). Dissociation and the fragmentary nature of traumatic memories: Overview and exploratory study. *Journal of Traumatic Stress,* **8,** 505–525.

Wegner, D. M. (1994). Ironic processes of mental control. *Psychological Review,* **101**(1), 34–52.

(Received April 16, 1999)
(Revision received December 28, 1999)