THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY, 2013
Vol. 66, No. 7, 1368–1398, http://dx.doi.org/10.1080/17470218.2012.743572

Suppression in retrieval practice, part-set cueing, and negative priming memory: The hydrogen model

Gabriel A. Radvansky and Andrea K. Tamplin

Department of Psychology, University of Notre Dame, Notre Dame, IN, USA

A number of phenomena in memory have been explained using appeals to active suppression processes, including retrieval practice, part-set cueing, and the negative priming that is observed with associative interference. However, more formal attempts to capture such processes have been absent. This paper outlines the hydrogen model of memory retrieval, which aims to be a simple model with the modest goal of trying to explore what influence suppression would have on memory retrieval. This model contains a single activation component and a single suppression component in which suppression comes into play only after retrieval interference has been detected. This model was created to explore the plausibility and viability of ideas about the operation of suppression during memory retrieval. For hydrogen, the degree of suppression recruited is proportional to the amount of interference experienced. Overall, the pattern of human data was captured by the suppression model.

Keywords: Suppression; Inhibition; Retrieval practice; Part-set cueing; Negative priming.

In cognitive psychology there is an overarching interest in attentional mechanisms of active suppression, in which the activation level of a mental representation is actively reduced, possibly below some baseline. This interest has been extended to memory retrieval with suggestions that suppression is available to a wide range of processes at many levels of analysis, from simple neural suppression (Carlson, 2004) to the suppression of complex autobiographical memories (Barnier, Hung, & Conway, 2004). The focus here is on the suppression of individual memory traces.

As a counter to these ideas, it has also been proposed that suppression is theoretically unnecessary (e.g., MacLeod, Dodd, Sheard, Wilson, & Bibi, 2003). Phenomena that have been attributed to active suppression can be accounted for by activation-only processes. For example, an inhibition *effect* may be due to a type of blocking interference arising out of an increase in activation of some memory traces at the cost of others. In this paper, to be clearer, we use the term *suppression* to refer to the cognitive mechanism and *inhibition* to refer to an observed effect.

There are differences in how people define suppression, and these differences impact how processes are interpreted. Here we delineate three types of suppression to clarify what our focus is.

Correspondence should be addressed to Gabriel A. Radvansky, Department of Psychology, University of Notre Dame, Notre Dame, IN 46556, USA. E-mail: gradvans@nd.edu

We would like to thank Sidney D'Mello, Melissa Lehman, and Jeroen Raaijmakers for their comments on earlier versions of this manuscript.

We then address motives for including suppression in memory retrieval, as well as some counterarguments. Following this exposition, we present the hydrogen model along with a description of the retrieval practice paradigm, part-set cueing, and the negative priming observed with the fan effect, along with the application of hydrogen to these effects.

ACTION REACTION

Here we derived an outline of three classes of suppression based on the role they play in cognition. The first is *structural suppression*. These processes are automatic, part-and-parcel processes of a memory structure or process, such as the inhibitory links in a parallel distributed network. This is structural suppression because the inhibitory links are part of the structure of the cognitive representation/process.

The second type is *metacognitive suppression*. This is suppression that is not a part of retrieval, per se, but is metacognitively imposed to avoid retrieving certain memories, as with directed forgetting (Bjork, 1970, 1989) in which a person metacognitively avoids recalling items that are designated as irrelevant in the current task. Thus, this is the suppression of the memory traces because of a deliberate, conscious segregation of memories as being relevant or irrelevant. Note that this may co-opt retrieval processes involved in retrieval suppression for this purpose.

The third type of suppression is *retrieval suppression* that is experienced via the act of retrieval itself (e.g., Healey, Campbell, Hasher, & Ossher, 2009). This type of suppression is present when multiple memory traces are activated during retrieval and compete with one another, thereby producing interference. When this interference is detected, suppression is used to remove the related but irrelevant traces from the retrieval process. Moreover, the amount of suppression that is used is proportional to the amount of interference experienced (e.g., Anderson, Bjork, & Bjork, 1994). The more interference there is, the more suppression there will be. If there are no competing memory traces, then no suppression is used. Hydrogen focuses on this third type of suppression as it is observed in episodic memory situations.

HYDROGEN

The hydrogen model presented here is a *simple* model that includes a single *activation* component and a single *suppression* component during retrieval, hence the name. The aim of this model is to go beyond verbal models of suppression (e.g., Anderson et al., 1994; Bäuml, 2002; Bäuml & Aslan, 2004; Radvansky, 1999b) to assess the viability of theories of memory that involve suppression and to provide a starting point for assessing claims about the operation of suppression during memory retrieval. It is not intended as a complete model of memory, such as more complex computation models of memory retrieval (e.g., Anderson, 1983; Hintzman, 1988; Norman, Newman, & Detre, 2007; Raaijmakers & Shiffrin, 1981) that are intended to cover a very wide range of memory phenomena—just a simple assessment of some basic principles. Hydrogen was created with a number of assumptions in mind. These are grouped into four categories. The first are assumptions about the nature of the memory representation. The second is the criterion used to determine retrieval. The third has to do with the activation of memory traces during retrieval, and the fourth with the suppression of competitors.

Representation

The model makes no strong theoretical assumptions about the form of long-term memory representations. Knowledge can be represented as concepts, propositions, mental models, mental images, and so on. These kinds of qualities are not important for the current questions being addressed. What is important is the idea that different units of information are treated as different traces.

Second, memory traces have an activation level prior to retrieval. In many cases, this level is at a normative baseline, above a minimum level. For a subset of traces, there are some that have higher or lower levels relative to the baseline because they were recently involved in processing, and their nonbaseline level reflects the results of being activated or suppressed. As a convention, for the simulations reported here, the range of activation levels is from 0 to 1, and the baseline level is .2. This level is selected because it is low enough to allow traces to be retrieved above some threshold, but high enough to allow traces to be suppressed below baseline.

Retrieval criterion

The second set of assumptions has to do with the criteria for designating a memory trace as retrieved. First, in keeping with the spirit of having a suppression mechanism in retrieval, rather than a threshold value, the criterion for the retrieval is the ability to discriminate a memory trace from its competitors. Thus, the retrieval criterion is that a trace's activation value exceeds its closest competitors by a certain amount. Here, more moderate values, such as .1 or .2, are more plausible given a range of activation values from 0 to 1.

Second, the criterion that people use is flexible, thereby serving as a source of bias. It can shift around depending on the circumstances. However, while this point is acknowledged, for all of the simulations reported here, it is fixed at .1. So, if all of the other memory traces in the retrieval set are at .2 or below, then an activation level of .3 would need to be reached. However, if at least one of them is higher, then the target trace would need to be at least .1 higher than that. For example, if one of the competitor traces had an activation value of .35, then the retrieved trace would need an activation value of at least .45. Thus, items are retrieved when they are sufficiently differentiated from the other traces in a retrieval set.

Memory activation

The third set of assumptions has to do with the retrieval activation. First, during the initial retrieval period, only traces that are related to a probe are activated (e.g., Neely, 1977). This retrieval is done by a global matching process. Those traces that share features, elements, or components with those in the memory probe are activated. Second, a memory probe can be broken down into components. The nature of these components would be specified by a theory using the principles outlined by hydrogen. Third, the traces that are activated by a probe are not activated to an equivalent degree. Similar to other theories, those traces that provide a closer match to the memory probe receive greater facilitation (e.g., Hintzman, 1988). So, facilitation is a function of the degree of probe–trace overlap, with the traces having more overlap being activated to a greater degree. Also, any traces that have been activated because of recent retrieval will be more likely to be activated, and thus this process serves as a means to block the access of its competitors.

We assume that a set amount of activation is available and that the total facilitation during retrieval sums to 1. During retrieval, this activation is divided among the components in the memory probe, which is then further divided among those traces that contain each of those components. This is the ratio rule that forms the basis of many models of memory retrieval (e.g., Anderson, 1983; Raaijmakers & Shiffrin, 1981). In sum, facilitation is distributed to the memory traces according to three principles: (a) the facilitation available is divided by the number of components (features) in a memory probe; (b) the amount of facilitation provided by each component is divided by the number of traces that have that component; and (c) the facilitation applied to each memory trace is the sum of the facilitation levels from each component.

Suppression

The final set of assumptions relates to the most interesting part of the model—namely, the operation of suppression. First, as noted earlier, during the initial retrieval period, a number of traces are activated as a function of their overlap with the memory probe. These traces compose the retrieval set. Suppression is not used until the retrieval set is defined, and it is clear that there is some interference to resolve (e.g., Bajo, Gómez-Ariza, Fernandez, & Marful, 2006; Healey et al., 2009). Again, interference occurs when there are multiple memory traces activated by a probe. The time between the onset of activation and the onset of suppression can potentially vary. However, in the simulations reported here, it is a constant (10 cycles of the model).

Second, most of the traces in the retrieval set are irrelevant to the task and so serve as sources of interference. As such, those traces should be suppressed. The rest of the traces are outside of the retrieval set. Interference is defined in terms of the amount of competition among traces in the retrieval set. The more similar they are to one another (in terms of their activation level) after the initial retrieval phase, the more interference there is.

Third, the application of suppression is not uniform. Instead, it is proportional to the amount of interference that is experienced. This is essentially extending the ratio rule used with activation to suppression. When there is a large amount of interference, there will be a large amount of suppression. In contrast, when there is a smaller amount of interference, there will be less suppression (e.g., Anderson et al., 1994). The level of interference is defined here as the compliment of the activation level difference among all of the traces in the retrieval set. Based on the amount of interference present, the necessary level of suppression can be set.

Fourth, after the degree of interference is determined, the change in trace activation levels then reflects the sum of the activation and suppression influences. The degree to which a memory trace has been activated influences the degree to which other traces are suppressed. Thus, both activation and suppression are influenced by the number and type of components present in the memory probe and the traces in memory. Furthermore, the amount of suppression from each trace is divided by the number of competitors it has, so that it sums to whatever the suppression level may be. As a result, the overall suppression process is a function of the degree to which different traces have been activated and the size of the retrieval set.

THE PROGRAM

This section defines some peripheral assumptions that were pragmatic choices made for the computer program. In terms of representation, each memory trace was an array of text-based features. Array positions had no importance. The program did not take into account any relations among features, although such factors can influence memory under other circumstances, such as serial order and meaningful associations. Moreover, although, theoretically, retrieval is continuous, it is difficult to actually implement it. So, the hydrogen program advanced in small time steps, updating everything as a consequence of what had occurred prior. So, each trace accumulated activation and suppression at a constant rate, which was the rate of change in activation levels per time step. This rate can be any value between 1 and 0, with smaller values providing more fine-grained estimates, closer to a continuous process. For all of the reported simulations, this rate was fixed at .005 (i.e., any facilitation or suppression effects were multiplied by .005). For example, it would take 20 time steps for a memory trace, without any sources of interference, and perfectly matching a memory probe, to be increased by .1.

RECOGNITION

In this section, we outline the process that the model goes through during recognition. Prior to a memory probe, each trace is at some baseline level of activation. Again, for all of the examples here, we assume a baseline of .2. The actual activation level of a trace at the beginning of a trial could vary depending on recent experience.

Retrieval set and division of facilitation

The model first identifies those traces that share components with the probe to create the retrieval set. The total facilitation available is set to 1. For simplicity, this facilitation is divided equally among the number of probe components, following the ratio rule. Furthermore, the facilitation allocated to a given memory for a given component is divided by the number of traces that share that element. The formulae for this and other processes are in Appendix A.

Activation, suppression, and decay

During the initial retrieval phase, all of the traces in a retrieval set increase in activation according to the amount of facilitation they receive. The activation level of a given trace is a function of its prior activation level plus any additional activation received. After a period of time, the level of interference is determined in terms of the difference in the activation levels of those traces in the retrieval set.

Once the level of interference is established, then the amount of suppression allocated is determined. If there is only one trace in the set, then no suppression is used. Otherwise, the amount used is a function of the degree of interference. The more interference there is, the more suppression there will be. Furthermore, suppression is divided in a manner similar to the division of facilitation based on the number of traces and probe components involved. Suppression comes from all of the memory traces receiving facilitation based on how much facilitation each trace has received, modified by the level of interference. For example, assume for the sake of argument that a memory trace has a facilitation level of .5, the interference level is .75, and there are two competitor traces. Thus, there would be a suppression level of .375 (.5 times .75) available to this trace, which would then allocate .188 to each of the two competitors. Thus, the level of suppression a trace is *receiving* is the sum of the suppression inputs from its competitors.

After the suppression levels are set, both the facilitation and suppression continue to alter a trace's activation level. Retrieval continues until a trace reaches the criteria level of difference in activation level relative to the other traces in the set and has exceeded the criterion level relative to the baseline level. Traces that are not part of the current retrieval set but were recently activated or suppressed decay back to baseline. Hydrogen uses a negatively accelerating power function for this decay process based on how far a given trace is from baseline at that moment.

Retrieval and negative responses

Once a trace reaches or exceeds criterion it is retrieved. At this point, hydrogen checks whether the components in the probe are in the retrieved trace. If not, perhaps because a distractor trace is at a higher level of activation prior to retrieval, then that trace is removed from the retrieval set and begins decaying. Retrieval continues with the reduced retrieval set.

When a lure probe is involved, to keep from indefinitely extending retrieval there is a stopping rule. Hydrogen stops searching when two, somewhat arbitrary, conditions are satisfied. First, for the simulations reported here, retrieval progresses until it is at least three times as long as it would take a single trace (i.e., no competitors) to be retrieved. Second, to allow the model to continue a bit further if it is getting close, there is a positively accelerating function based on where retrieval starts. After this adjusted time period, if the top activation level in the retrieval set is still not sufficiently beyond the criterion, then retrieval stops.

Occasionally retrieval stops when the target is in memory, and the model incorrectly produces a negative response (a miss). There is also a probability that the model will make misses and false alarms as a function of the fluency of the retrieval (as indexed by the amount of time needed) with longer times increasing the error probability, and the amount of interference experienced on a trial, with greater interference levels increasing the probability of an error.

RECALL

In this section, we outline the process that the model goes through during recall. Again, prior to retrieval, each trace is at some level of activation, with a baseline of .2.

Retrieval set, sampling, and division of facilitation

For recall, all of the traces are part of the retrieval set. The model samples from this set, pseudorandomly influenced by the current activation levels, with traces at a higher level of activation being more likely to be sampled than traces at lower levels. Again, this is essentially the ratio rule, a principle that is used in many other memory models. Traces that were suppressed have a below-chance level of being selected. Once a trace reaches or exceeds criterion, in terms of both the other traces and the baseline level of activation, it is retrieved. Once a trace is reported, it is removed from the retrieval set. Hydrogen only excludes the most recently recalled item in the retrieval set. Other previously recalled items could be resampled at that point.

During recall, like recognition, all of the traces in a retrieval set increase in activation according to the amount of facilitation they receive. This process continues until a trace reaches criterion relative to all of the traces and has exceeded the criterion level relative to the baseline level. Traces that are not part of the retrieval set decay back to baseline. The activation of related traces increases the possibility that they will be recalled later.

After each trace is retrieved, there is a check to assess whether it has been recalled before. If so, then no response is made, and that trace begins to decay. Recall stops after there are a number of unsuccessful attempts to generate new information. For the simulations reported here, as a convention, this stopping point was when the number of unsuccessful recall attempts was half the number of items in the set. Potentially this stopping point could vary to capture different individuals' retrieval criteria.

SIMULATIONS

We applied hydrogen to three retrieval phenomena that have had theoretical arguments made for the operation of suppression. These are the retrieval practice, part-set cueing, and fan effects. In addition, we also did an assessment of the impact of different levels of interference in the retrieval practice and fan effect paradigms. We consider each of these in turn.

Retrieval practice

The retrieval practice paradigm produces an effect of reduced recall of previously studied items when members of the same category are repeatedly retrieved prior to final recall. This effect is usually elicited using a process involving three phases. First, people are presented with category–item pairs in the study phase. In the second phase, half of some of the categories are recalled by the prompt of category name, plus the first two letters of an item. People fill in the blank with the rest of the letters from the words they learned in the learning phase.

Finally, after a delay, in the third phase (test phase) people try to recall all the category–item pairs from the study phase. The items that were practised in the study phase are called RP+ items, whereas the items that were not practised during retrieval practice, but were from the same categories, are called RP– items. Finally, items from categories that were only studied and not practised at all are called NRP items. For example, suppose a person was given a list that had the categories *tool* and *tree*. If *tool* was the to-be-practised category, and items like tool–hammer, tool–pliers, and tool–saw were practised, but tool–screwdriver, tool–drill, and tool–vice were not, the first three would be RP+ items, and the second three would be RP– items. All of the items from the category *tree* (e.g., tree–oak) would be NRP items.

Retrieval practice frequently produces a pattern of data in which the RP+ words are more often recalled during the final test than NRP words, which in turn are recalled more often than the RP– words. Activation-only accounts (e.g., Jakab & Raaijmakers, 2009; William & Zacks, 2001) suggest that there is increased activation for memory traces related to the cue. As such, these items are more likely to be recalled. The decline in memory for the RP– items is because these items are at lower activation levels and so are less likely to be selected during recall.

In contrast, active suppression views suggest that the related but irrelevant traces are repeatedly suppressed during practice (Anderson et al., 1994). Thus, their activation levels are below baseline, causing the RP– items to be recalled at a lower rate than NRP items. In addition to the retrieval practice effect itself, there are other sources of evidence in support of a suppression view. For instance, there is evidence of such suppression across a wide variety of memory tasks, including category-stem cued recall (e.g., Anderson et al., 1994), recognition (e.g., Hicks & Starns, 2004), and even a variety of implicit memory tasks (Camp, Pecher, & Schmidt, 2005; Perfect, Moulin, Conway, & Perry, 2002). Moreover, the presence of cue-independent suppression—that is, the inhibition of items that are semantically related, but not paired with the practice cue—has also been used as an indicator of active suppression in producing the retrieval practice effect (e.g., Anderson & Spellman, 1995; Camp et al., 2005).

The suppression hypothesis is also supported by neuroimaging data. For example, electroencephalography (EEG) data has shown that there is greater mental activity when there are interfering memories that would be candidates for suppression (Johansson, Aslan, Bäuml, Gäbel, & Mecklinger, 2007; Staudigl, Hanslmayr, & Bäuml, 2010). Also, functional magnetic resonance imaging (fMRI) data have shown that as retrieval practice progressed, and memories were presumably suppressed, there was less activity in the prefrontal cortex (Kuhl, Dudukovic, Kahn, & Wagner, 2007; Wimber, Rutscuhmann, Greenlee, & Bäuml, 2009). This is consistent with the idea that there was less need for cognitive control under those circumstances because there would now be less interference.

The retrieval practice task has previously been simulated using a neural network model that uses hippocampal (episodic memory) and cortical (item, associative, and contextual) sources to simulate retrieval (Norman et al., 2007). In essence, the Norman et al. model uses a process of unlearning to weaken connections in the network to bring about the retrieval practice effect. It is not clear that this effect involves an unlearning process given that these effects can disappear with time (MacLeod & Macrae, 2001). The model would require a process of restrengthening weakened connections in the absence of new learning. Norman et al. speculate that a rapid eye movement (REM) sleep process could serve this purpose. However, there is no evidence to support this idea. Our aim here is not to evaluate Norman et al.'s model, which we think has a number of merits. However, given that this model does not allow one to easily assess the effects of an activation-only process as a comparison, and it has not been extended to other retrieval induced forgetting phenomena, there is value to pursuing the hydrogen simulations reported here.

We compared the suppression and activation-only accounts of retrieval practice using hydrogen to model retrieval when suppression is and is not involved. This simulation was done twice, once with suppression and once without. Note that for the activation-only simulations, retrieval practice effects might emerge because the prior retrieval of some items increases their activation level and, thereby, their probability of being recalled later, functionally blocking access to the nonpractised items. Thus, hydrogen with suppression turned off captures the basic processes argued to account for the retrieval practice effect in activation-only accounts. Forty-eight simulation subjects were generated using materials derived from the process described by Anderson and Spellman (1995), Anderson and Bell (2001), and Anderson et al. (1994).

Recall

In the standard paradigm, simulated subjects were given lists of 60 words in 10 categories to establish the information in long-term memory. During practice, half of the words in half of the categories were practised. Following this was free recall retrieval. The model went through two separate retrieval conditions. These were: (a) with suppression on, and (b) with suppression off. For the suppression-on trials, the model suppressed competitors, whereas for the suppression-off trials, only the processes of activation and decay were operating. This is the procedure used for all of the simulations reported here.

The pattern of influence operating during retrieval practice on a given trial is shown in Figure 1, with the lines between the various components reflecting sources of facilitation and suppression. Here, each trace has a category name and the name of the exemplar that is in that category. To make this example easier to follow, we assume that there were four items in each category, rather than six as was done in the actual simulations. Because this cued recall procedure effectively targets a single memory trace, we greatly simplify things and treat this procedure as equivalent to recognition. The amount of facilitation each trace receives was divided by the number of exemplars, which is the same for all of the categories here. So, the amount of facilitation would be divided into two (.500), half for the category, and half for the item. The facilitation for the category would be divided by four because there are four memory traces related to the category. So, each memory trace would receive .125 of the facilitation, except for the practised trace, which would receive .625 of the facilitation.

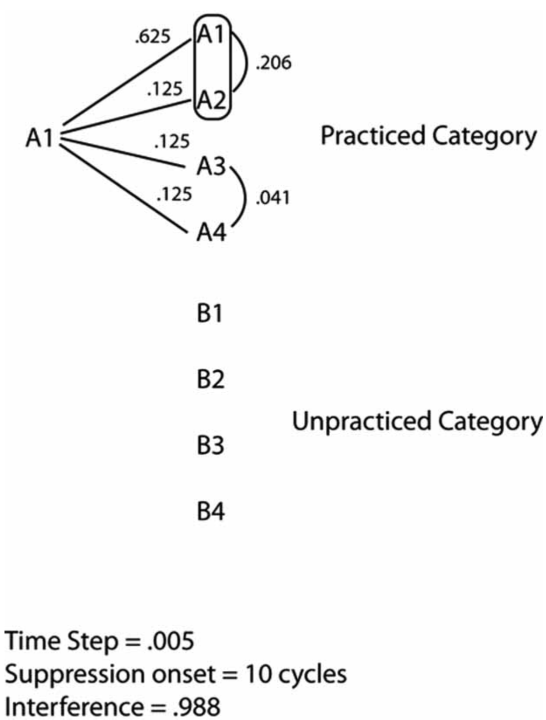

Figure 1. *Retrieval practice structure for a given cue (Item A1) on memory traces from Category A, Items 1–4, such as when a person is given the cue "Leather", and Items 1, 2, 3, and 4 are "gloves", "purse", "wallet", and "belt". A facilitation value of .625 goes to the memory trace for Item A1, whereas .125 facilitation goes to the rest of the items in that category. There is .206 suppression going from Trace A1 to the others (not all suppressions are shown) and .041 suppression from the competitor traces to the others (not all suppressions are shown).*

Because there are multiple traces in the retrieval set during the practice phase, there is interference. Interference levels are calculated as described earlier. Assuming all activation levels are equal at the start, the interference level would be .988 after 10 cycles with a time step of .005. After it has been established that there is interference, then suppression starts. To determine the suppression level coming from each memory trace, the facilitation level coming in is multiplied by this interference value. Suppression is then divided among the competitor traces. So, for the target trace, the suppression level is .625 × .988 (.618) divided by the three other traces it is applied to (.206). In comparison, for the distractor traces, the suppression level would be .125 × .988 (.124) divided by the three other traces it is applied to (.041).

The amount of suppression a trace receives is the sum of the suppressive influences. The target trace gets a .123 suppressive influence. When this suppression is combined with the facilitation level (.625 – .123), the overall facilitation level is reduced (.502). In comparison, the competitor traces receive a .288 suppressive influence; when this is combined with the facilitation level (.125 – .288), the overall activation level change is negative (–.163). This is active suppression. If activation levels prior to the onset of a cue are different (from previous trials), then the complexity of the computations increases. Specifically, this change would alter the levels of interference experienced and, consequently, the division of suppression.

The retrieval practice data are shown in Figure 2. Figure 2a shows the weighted (based on number of subjects) human data, listed in Appendix B. These data are from studies using word lists in which the lists were categorized, people practised half of the items from half of the categories, and the memory test was a recall test. Figure 2b is hydrogen with suppression, and Figure 2c is hydrogen without suppression. There was no evidence of a retrieval practice effect when no suppression was involved. The pattern of data produced by hydrogen with inhibition closely resembles the human data. Specifically, for both the human and hydrogen with suppression data, performance is worse in the RP– condition. To show that suppression is an essential ingredient of hydrogen to simulate human data, Figure 2c shows that when suppression is not operating, while there is superior performance for those items that were practised, performance was similar in the RP– and NRP conditions.

One further point to note is that if performance of the model on the RP+ and RP– items is submitted to a correlation analysis, the result is that there is no significant correlation, $r = -.03$. Thus, as it is produced by hydrogen, the retrieval practice effect (inhibition) is different from the processes that produce the superior memory for the practised items (facilitation). This is in line with other research that has similarly found an absence of such a relationship (e.g., Aslan & Bäuml, 2011; Staudigl et al., 2010), including EEG measures (Hanslmayr, Staudigl, Aslan, & Bäuml, 2010), which is inconsistent with activation-only accounts.

Recognition

While most studies using the retrieval practice paradigm involve a final recall test, as was used in the previous section, there has been some suggestion that recognition would be a more direct means of assessing the impact of suppression on performance (Aslan & Bäuml, 2010). For both suppression and activation-only models, people should respond more accurately to the RP+ items than the NRP items because the RP+ items were recently practised, thereby increasing their activation level, rendering them highly accessible.

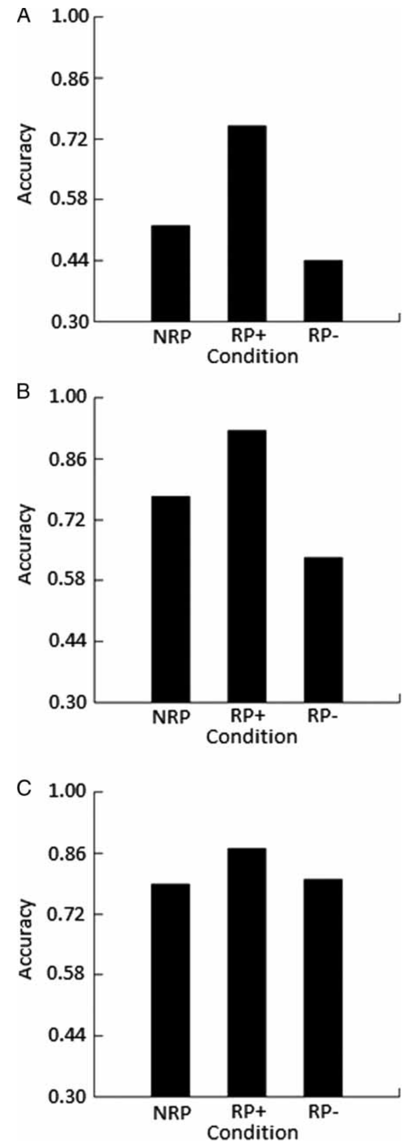

Figure 2. *Results of the retrieval practice experiments with recall and simulations. A is the human data average across many studies, B is hydrogen with suppression operating, and C is hydrogen without suppression operating. RP+ =items practised in the study phase; RP– = items not practised during retrieval practice, but from the same categories; NRP = items from categories that were only studied and not practised at all.*

Beyond this prediction, a suppression model predicts that performance on the RP– items will be more error prone than that for the NRP items.

This prediction is because those items were suppressed below baseline, and it will be more difficult to move them into a retrieved state. Conversely, an activation-only model predicts that performance on the RP– items will be less error prone than the NRP items. This prediction is because the continuous practice of the RP+ items will implicitly prime the RP– items, and, as a result, people will respond more accurately to these items on a recognition test.

To address this issue, the hydrogen model was used as before; however, rather than recall after retrieval practice, there was a recognition test. The recognition data are shown in Figure 3. Figure 3a shows the averaged human data. Figure 3b is hydrogen with suppression, and Figure 3c is hydrogen without suppression. Again, there is no evidence of a retrieval practice effect when no suppression is involved, but the pattern of data produced by hydrogen with inhibition more closely resembles the human data. Specifically, for both the human and hydrogen with suppression data, responses are less accurate in the RP– condition than in the NRP condition. In contrast, when suppression is not operating, performance is similar in these two conditions.

One question that might arise at this point is why might a basic activation-only mechanism not produce the retrieval practice effect? The logic of an activation-only account is that by repeatedly retrieving a subset of a category, those particular items are increasing in their baseline activation level and so are more likely to be sampled during later retrieval. As a consequence, those items that were not practised should be less likely to be sampled, so that, proportionally, the probability of one of those items being sampled would be less than that of an item from a nonpractised category (the retrieval practice effect). We agree with this characterization up to this point, and in the hydrogen simulations this is occurring.

However, there is something else simultaneously occurring for hydrogen that is not factored in by the prior activation-only accounts—and that is the influence of (positive) priming. Every time one of the items of a category is subject to retrieval practice, all of the other items in the category are primed, thereby boosting their base activation level and boosting the retrieval level overall. As a result, in the case of these simulations, the negative and positive consequences more or less cancel one another out. One could produce the repeated

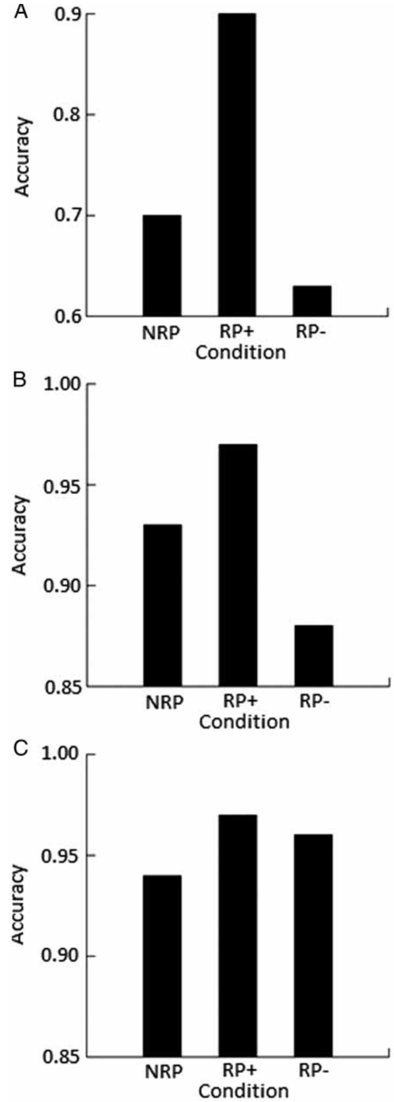

Figure 3. *Results of the retrieval practice experiments with recognition and simulations. A is the human data average across several studies, B is hydrogen with suppression operating, and C is hydrogen without suppression operating. RP+ =items practised in the study phase; RP– = items not practised during retrieval practice, but from the same categories; NRP = items from categories that were only studied and not practised at all.*

practice effect using an activation-only version of the model by eliminating priming. However, that would go against decades of cognitive evidence to the contrary, given how strong priming effects have been shown to be (e.g., the Deese–Roediger–McDermott, DRM, false memory effect), and would go against the ratio-rule of the division of activation that is the basis of many cognitive models.

So, we acknowledge that an activation-only account could potentially produce the repeated practice effect. However, our simulation results raise an important issue: How do blocking and priming processes simultaneously influence one another in this experimental paradigm? This is not an issue for models that involve suppression because the positive priming is directly counteracted by the suppressive influence. For activation-only accounts, when these two basic, and commonly held, principles are implemented—namely, that (a) retrieval practice increases the activation level of those items, and (b) there is a spread of activation to related memory traces—then the repeated practice effect may not emerge. A more complex activation-only model would need to be invoked to explore different theoretical possibilities.

Part-set cueing

The part-set cueing effect is a decrease in the recall of items from a set when a person is provided with a subset of those items, relative to when they are told nothing (Slamecka, 1968). A number of explanations have been provided for the part-set cueing effect. One idea is that when people are presented with part-set cues, this disrupts their retrieval plan and imposes a different organization of the material (e.g., Basden & Basden, 1995; Basden, Basden, & Galloway, 1977). This inconsistency disrupts retrieval, and so there is a performance cost (Slamecka, 1968).

Another explanation involves active suppression (e.g., Aslan & Bäuml, 2007; Aslan, Bäuml, & Grundgeiger, 2007; Bäuml & Aslan, 2004, 2006). Specifically, giving part-set cues causes the implicit retrieval and recognition of that information. At this time, the related memory traces are treated as competitors and are suppressed. It is this idea that is tested with hydrogen.

In addition to the part-set cueing effect, itself, there are a number of behavioural indicators to support the idea that active suppression is involved. For example, part-set cueing results in worse recognition performance. This is expected according to activation-only accounts that emphasize the organizational structure of recall (e.g., Oswald, Serra, & Krishna, 2005). There is also evidence of inhibition when there is independent-probe testing (Aslan et al., 2007), consistent with what has been found with retrieval practice. The presence of inhibition effects in recognition, as well as recall, is also consistent with findings using the retrieval practice paradigm, supporting the idea that a common mechanism underlies them both. Finally, the active suppression account has also been supported by research using fMRI recordings (Crescentini, Shallice, Del Missier, & Macaluso, 2010). In part-set cueing trials, as compared to uncued controls, there is greater activation in the left prefrontal and right dorsolateral prefrontal cortices, areas that are associated with the regulation of interference and suppression.

There are some activation-only accounts of part-set cueing (e.g., Raaijmakers & Shiffrin, 1981; Rundus, 1973). However, these accounts depend on the idea that people have structured the information in memory and that the retrieval cues disrupt the retrieval plan. So, the effect in these accounts is due to the organization of the items during encoding and storage, not by the retrieval process per se. While such accounts are plausible on their own, and some aspect of a mismatch between the mental organization and the retrieval cues does affect performance, it cannot explain the retrieval practice and fan effects also studied here.

Recall

For the part-set cueing task, each hypothetical subject learned a set of 24 category–exemplar pairs (e.g., hammer–tool), with four categories and six exemplars per category. The model then went through four memory retrieval conditions for each subject. These were: (a) uncued with suppression on, (b) cued with suppression on, (c) uncued with suppression off, and (d) cued with suppression off. For the uncued recalls, the model simply recalled as much as possible. In contrast, for the cued trials, the model was given three exemplar–category pairs from each category as the part-set cues. These were treated like recognition trials so suppression could operate. Essentially, these recall cues are precise enough to target a single memory trace as would occur in recognition. The categories were presented in a blocked fashion.

The pattern of influence operating at retrieval for a given category on a given trial is shown in Figure 4. The lines between components reflect sources of facilitation and suppression. In these experiments, each trace has a category name and the names of the exemplars that are in that category. The amount of facilitation each trace receives was divided by the number of exemplars, which is the same for all of the categories here. So, the amount of facilitation here is divided into two (.500), half for the category, and half for the item (i.e., red gets .500, and strawberry gets .500 activation). The facilitation for the category would be divided by six because there are six memory traces. So, each memory trace would receive .083 of the facilitation, except for the cued trace, which would receive .583 of the facilitation.

Because multiple traces are competing during retrieval, assuming that all activation levels are equal at the start, the interference level is .992 after 10 cycles of the model with a time step of .005. After it has been established that there is interference, then suppression starts. To determine the suppression level coming from each trace, the facilitation level coming in is multiplied by this interference value. This suppression is then divided among the competitor traces. So, for the target trace, the suppression level is .583 × .992 (.578) divided by the five other traces it is applied to (.116). In comparison, for the distractor traces, the suppression level is .083 × .992 (.082) divided by the five other traces it is applied to (.016). The amount of suppression a given trace is receiving is the sum of the suppressive influences.

So, the target trace receives a .082 suppressive influence. When this influence is combined with the facilitation level (.583 – .082), the overall facilitation level is reduced (.501). Thus, this trace is facilitated overall compared to the others and eventually is recalled. In comparison, the competitor traces receive a .180 suppressive influence; when this is combined with the facilitation level (.083 – .180), the overall facilitation level is negative (–.097) (suppression below baseline). If the activation levels prior to the onset of a cue are different (because of previous trials), then the complexity of the computations increases accordingly.

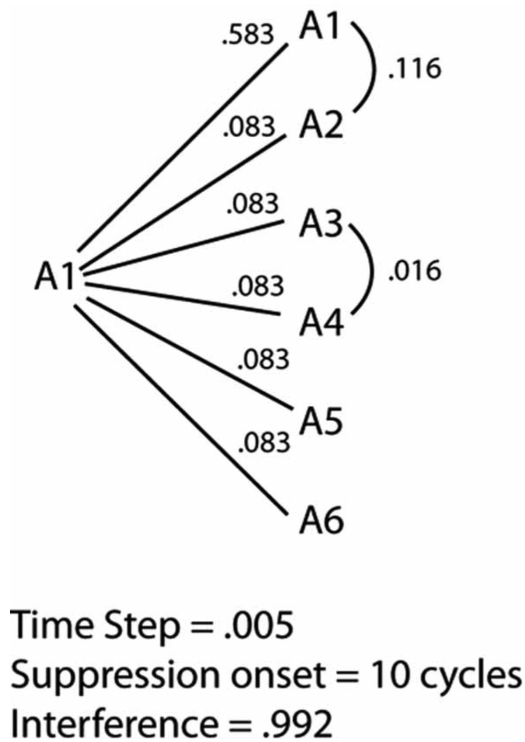

Figure 4. *Part-set cueing structure for a given cue (Item A1) on memory traces from Category A, such as "tool", and Items 1–6, such as "hammer", "screwdriver", "saw", "wrench", "pliers", and "level". The amount of facilitation for the A1 memory trace is .583, whereas .083 facilitation goes to the rest of the traces. There is .116 suppression going from Trace A1 to the others (not all suppressions are shown) and .016 suppression from the competitor traces to the others (not all suppressions are shown).*

The part-set cueing recall data are shown in Figure 5. Figure 5a shows the weighted, averaged human data. These data are listed in Appendix C. These data are from studies in which the lists

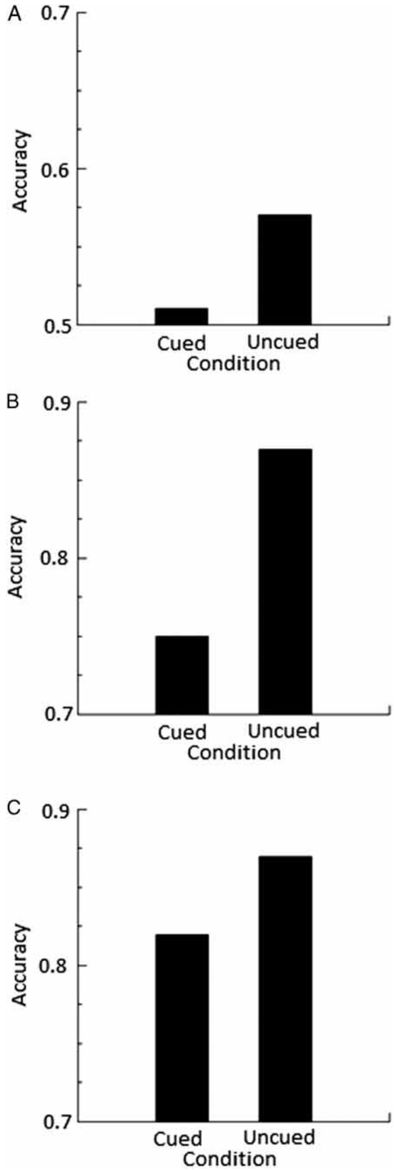

Figure 5. *Results of the standard part-set cueing recall experiments and simulations. A is the human data averaged across several studies, B is hydrogen with suppression operating, and C is hydrogen without suppression operating.*

were categorized, the part-set cueing was done within subjects, and the memory test was free recall. Figure 5b is hydrogen with suppression, and Figure 5c is hydrogen without suppression. While the pattern of data both with and without suppression produced a part-set cueing effect, the effect was much larger when suppression was involved. Specifically, performance was worse when people are given part of the list as a cue than if they are simply told to recall the list. Thus, it is possible for an activation-only model to produce a part-set cueing effect, just as has been shown before (e.g., Raaijmakers & Shiffrin, 1981; Rundus, 1973).

Recognition

For the recognition test data, everything was the same as with recall except that recognition trials were given following the part-set cues. The recognition data are shown in Figure 6. Figure 6a shows the weighted averaged human data. These data are listed in Appendix C. These data are from studies in which the lists were categorized, the part-set cueing was done between subjects, and the memory test was item recognition and the dependent measure was accuracy. Figure 6b is hydrogen with suppression, and Figure 6c is hydrogen without suppression. As can be seen, the pattern of data with suppression produced a part-set cueing effect, like the human data. However, this effect was reversed when suppression was not involved. Thus, while it is possible for an activation-only model to produce a part-set cueing effect when modelling recall data, we failed to observe such a pattern when a recognition test was used, although this is the pattern that people show.

Fan effect and negative priming

A fan effect is an increase in retrieval time with an increase in the number of associations with a concept (Anderson, 1974). In other words, memory retrieval is compromised when there are related but irrelevant memory traces. Radvansky (1999b; Radvansky, Zacks, & Hasher, 2005) assessed the role of suppression in a fan effect paradigm. People first memorized lists of sentences about objects and locations, such as “The potted palm is in the hotel”. Across the study list, there were 1 to 3 associations with each object and

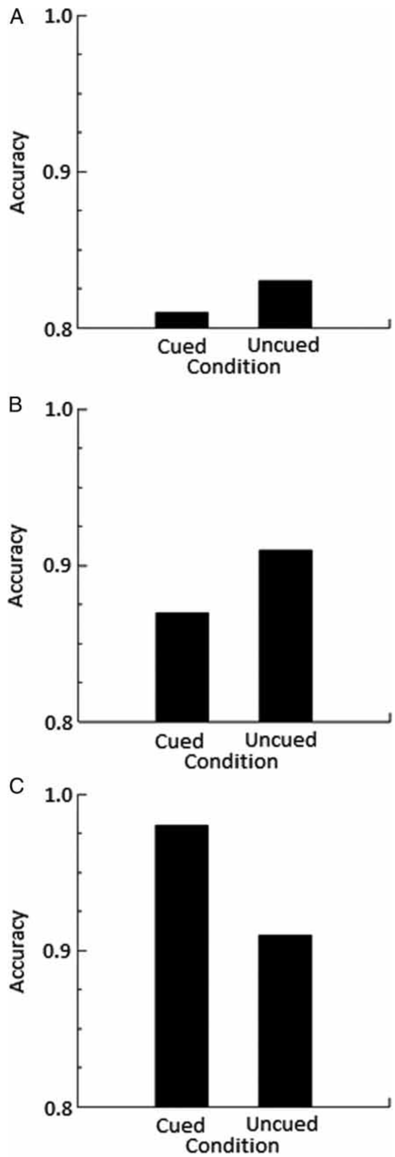

Figure 6. *Results of the standard part set cueing recognition experiments and simulations. A is the human data averaged across several studies, B is hydrogen with suppression operating, and C is hydrogen without suppression operating.*

location. Consistent with previous research, a fan effect was observed when there were multiple associations with a given object, but not when there were multiple associations with a given location (Radvansky, 1998, 2005, 2009; Radvansky & Copeland, 2006a; Radvansky, Spieler, & Zacks, 1993; Radvansky & Zacks, 1991; Radvansky, Zacks, & Hasher, 1996).

The interpretation of this pattern of data is that people are using situation models of the information in the study sentences. When a set of sentences refers to a single object in several locations, because each of these items is likely to be interpreted as referring to a separate situation, several models are created. Later, during recognition, because they share components with the memory probe, these models then interfere with one another during retrieval. The more irrelevant models there are, the greater the interference. In contrast, when a set of sentences refers to multiple objects in a single location, because these items can readily be interpreted as referring to a single situation, this information can be integrated into a single model. As such, during retrieval, there are no related but irrelevant models to produce interference.

The issue raised by Radvansky (1999b) was whether suppression operates when interference is observed. To explore this idea, a priming paradigm was set up in the context of a recognition test. A given target item was involved in both the experimental and the control conditions. In the *experimental* condition the probe that served as a prime (e.g., "The potted palm is in the hotel") contained an object that was in at least one other location in addition to the one in the probe itself, which results in interference from the related but irrelevant situation models. Following the prime, a target probe was presented (e.g., "The potted palm is in the museum"). This target required the retrieval of one of the related but irrelevant models from the previous trial because it contained the object from the prime trial. In the *control* condition the target was the same as that in the experimental condition, but the probe was unrelated to the target (see also Nievas & Mari-Beffa, 2002, for a similar paradigm involving irrelevant word meanings).

The relation between the fan and negative priming effects is shown in Figure 7. In the top of the figure, on trial *t*, a person is probed with an item that is about Object A in Location 1.

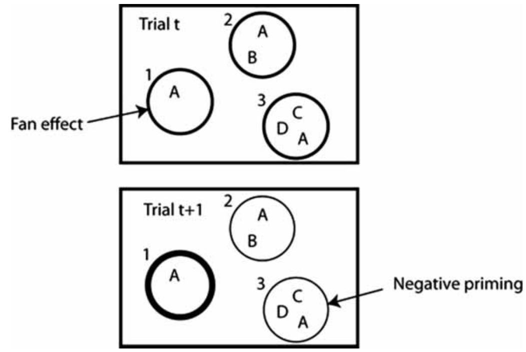

Figure 7. *Illustration of the sequence of trials during a recognition task that produces a fan effect on trial* t *and a negative priming effect on trial* t + 1*. Note that each circle would correspond to a situation model and the letters to individual objects in each event.*

Here, Object A is in three locations, and the interference that results produces the fan effect on trial *t*, with people responding slower than if Object A were in only one or two locations. Then on trial $t + 1$, people are given a probe that taps one of the traces that produced the interference on trial *t*. For example, people would be given a probe about Object C in Location 3. Although the object and location names were not mentioned in the previous probe, because Location 3 also has Object A, it was a source of interference on trial *t*. So, during retrieval, it was suppressed to a lower level of activation. Thus, response time on trial $t + 1$ is longer than it would otherwise have been. This is termed a negative priming effect because it is conceptually analogous to the distractor suppression effect reported in perception studies (e.g., Tipper, 1985). The focus here is not on the fan effect itself, but on the difference between the experimental and control trials that follow trials that produce the fan effect. This is because these are the trials that are diagnostic concerning the operation of suppression during the retrieval as a means of resolving interference.

Note that there is an account for the patterns of data modelled here using Adaptive Control of Thought – Rational (ACT-R; Anderson & Reder, 1999; Sohn, Anderson, Reder, & Goode, 2004); however, this account requires several implausible assumptions (Radvansky, 1999a, 2005), such as the requirement for concepts that have been learned as being associated, to be represented as being dissociated. More to the point, the negative priming effect that follows the fan effect trials have never been accounted for by activation-only models in more than a decade since the first results were reported.

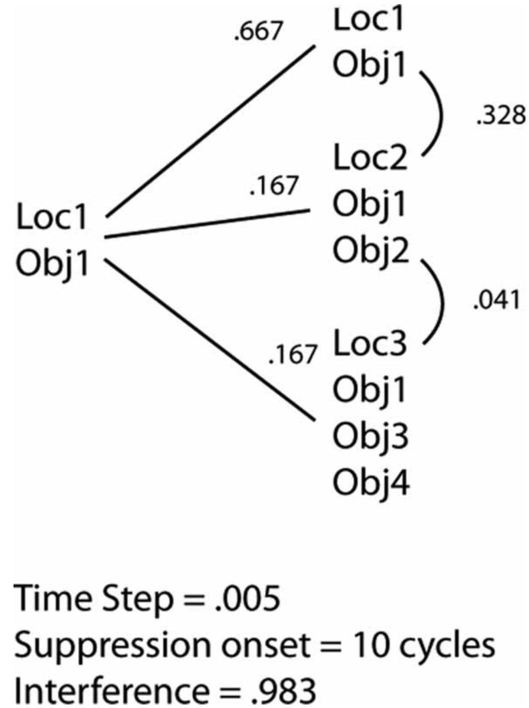

Figure 8. *Retrieval structure for a fan effect recognition trial in which the probe is Object 1 (Obj1) in Location 1 (Loc1), such as "potted palm" and "hotel", and there are two other locations in which the object was positioned, such as "airport" and "barber shop". The target situation model receives .667 facilitation, whereas .167 facilitation goes to the rest. There is .328 suppression going from Trace A1 to the others (not all suppressions are shown) and .041 suppression from the competitor traces to the others (not all suppressions are shown).*

To better illustrate how hydrogen captures retrieval in this paradigm, the process is shown in Figure 8. In these experiments, each trace has an object and a location. Given that the total amount of facilitation sums to 1, the amount of facilitation each trace receives is divided into two (.500), half for the location, and half for the object. Moreover, the amount of facilitation is divided by the number of models sharing the object. The facilitation for the object would be divided by three because there are three situation models. So, the irrelevant model traces here would receive .167 of the facilitation, and the target trace would receive .667 of the facilitation.

Because there is competitive retrieval, assuming equal activation levels at the start, the interference level would be .983 after 10 cycles of hydrogen with a time step of .005. After interference has been established, suppression starts. To determine the amount of suppression coming from each trace, the facilitation coming in is multiplied by the interference value. Suppression is then divided among the competitors. So, for the target, the suppression level is .667 times .983 (.656) divided by the two other situation models it is applied to (.328). In comparison, for the competitors, the suppression level is .167 × .983 (.164) divided by the two other models traces it is applied to (.082).

The amount of suppression a given trace receives is the sum of the suppressive influences. So, the target receives a .164 suppressive influence. When this influence is combined with the facilitation level (.667 − .164), the overall facilitation level is reduced (.503). In comparison, the competitors receive a .410 suppressive influence. When this influence is combined with the facilitation level (.167 – .410), the overall facilitation level is negative (–.243). As such, if one of these models is probed for on the next trial, it will take longer to retrieve it.

To run the simulations, the stimulus files for Radvansky (1999b, Experiment 2) were used. For the recognition test, each simulation subject got 216 trials. The studied probes were items from the study list. There were 6 trials for each studied item. The nonstudied probes were recombinations of objects and locations from within the same cell of the design. For the priming trials, the prime probe referred to an object that was also in the target location (which was a competitor on the prime trial). On the target trials, the probe referred to a competitor location that held the previous object. For the control priming trials, the target was the same, but the prime was from the same cell of the design as the experimental primes, so it had the same number of associations with the object and location concepts, but was not related to the target location. Because the experimental target trials are argued to not involve suppression, they are not included in the analysis of the fan effect per se. The fan effect occurs because of competition among memory traces, and this competition is resolved in part using suppression. The negative priming effect is a result of the recent suppression of related but irrelevant memory traces.

The fan effect data are shown in Figure 9. Figure 9a shows the weighted averaged human data, listed in Appendix D. Figure 9b is hydrogen with suppression, and Figure 9c is hydrogen without suppression. As can be seen, the differential fan effect observed in human data is present to some degree when suppression both is and is not involved.

Of more central interest is performance on the target trials. For the human data, we used the weighted average of Experiments 1 and 2 from Radvansky (1999b) and the younger adult data from Radvansky et al. (2005). These data along with the hydrogen simulations with and without suppression are shown in Figure 10. Here, hydrogen replicates the pattern of slower retrieval times in the experimental condition when there is active suppression. However, when there is only activation, there is essentially no difference, and, if anything, the opposite pattern emerges. Thus, it is plausible that suppression is involved to manage interference during retrieval in a fan effect paradigm as evidence by performance on the negative priming trials.

Level of interference and amount of suppression

An important point that was made earlier was that the amount of suppression involved during retrieval is a function of the degree of experienced interference. This is a point that has been made throughout the theoretical consideration of suppression as a retrieval mechanism (e.g., Anderson et al., 1994). However, the more careful consideration provided by hydrogen suggests that interference can have different influences on the inhibition effects that are observed. In this section we outline two types

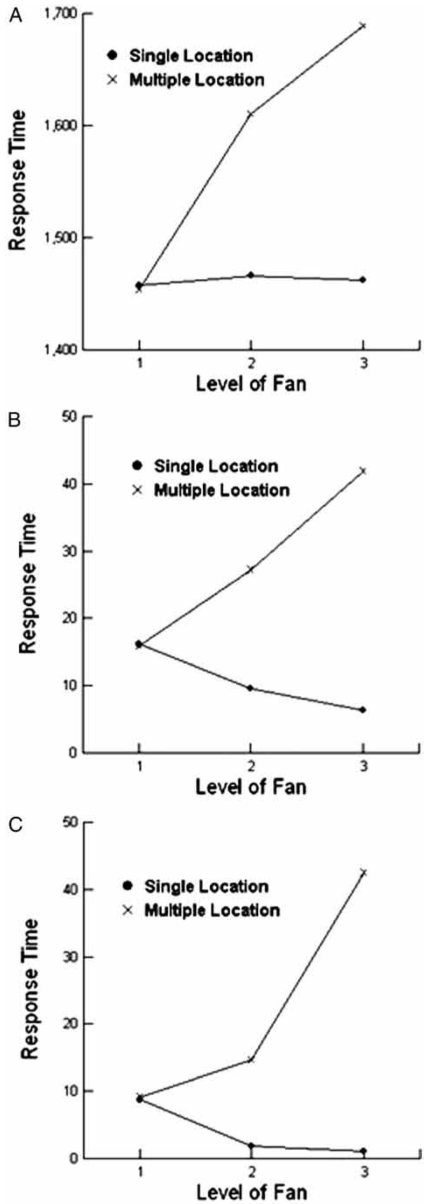

Figure 9. *Results of the fan effect experiments and simulations. A is the human data averaged across several studies, B is hydrogen with suppression operating, and C is hydrogen without suppression operating.*
Note: Units for A are in milliseconds, units for B and C are in cycles.

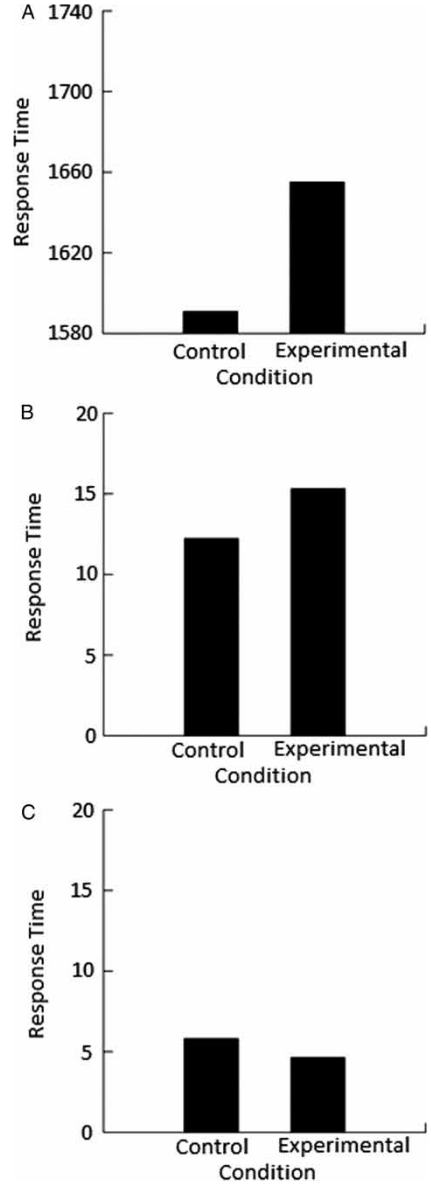

Figure 10. *Results from the negative priming component of fan effect experiments and simulations. A is the human data average across studies, B is hydrogen with suppression operating, and C is hydrogen without suppression operating.*
Note: Units for A are in milliseconds, units for B and C are in cycles.

of interference, provide the predictions from hydrogen, and then compare these with the observed patterns of data.

The first type of interference (cue strength interference) refers to how strongly associated an item is with a retrieval cue. This is the type of interference discussed by Anderson et al. (1994). In essence, memory traces that are weakly associated with a cue will produce less interference and, therefore, be exposed to less suppression and show a smaller inhibition effect. The second type of interference (competitor set size interference) refers to how many competitors there are during retrieval, as in a fan effect paradigm. So, more competitors generate more interference. However, as will be seen, this interference does not produce larger inhibition effects.

In terms of cue strength interference, memory traces vary in the degree to which they are activated by the cue. Competitors that are more strongly activated produce more interference than those that are more weakly activated. As a result, inhibition effects should be larger when the competitors are stronger category associates (Anderson et al., 1994). Such an outcome is possible using activation-only models only if, as noted by Jakab and Raaijmakers (2009), unreasonable assumptions are made. However, it has not been shown that an approach that includes active suppression produces such an outcome. To derive predictions for this using hydrogen, we used the retrieval practice paradigm. To manipulate the degree of memory trace activation by a retrieval cue, the category portion of the cue included two elements rather than one. More strongly related traces (half of the items) had both of these elements, whereas the more weakly related traces had only one. Otherwise, the simulation was as described earlier. The results are summarized in Figure 11, which are graphed as the degree of retrieval suppression (RP– minus neutral) to make the size of the effects clearer. As can be seen, the inhibition effect was larger when the competitors were stronger (produced more interference) than when they were weaker. So, when a competitor was more strongly associated with a cue, it is initially activated to a greater degree, and it also takes longer for the target trace to be retrieved. Both of these factors combine to lead to more suppression being applied to the competitor traces, resulting in a larger retrieval practice inhibition effect.

The simulation data were compared with the pattern of inhibition effects from similar studies (Anderson et al., 1994; Williams & Zacks, 2001), which are also shown in Figure 11 using the

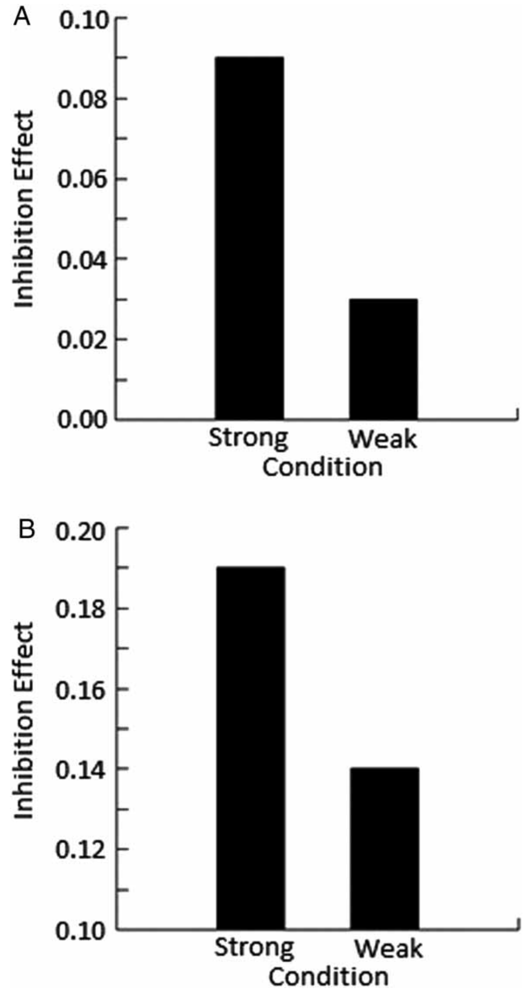

Figure 11. *Results from the cue strength retrieval practice experiments and simulation. A is the human data, and B is hydrogen with suppression operating.*

weighted average data across the studies. These means are reported in Appendix E. Note that while this effect was originally reported by Anderson et al. (1994), Williams and Zacks (2001) found the effect weak and inconsistent (along with overall lower levels of recall). Note also that an analogous strength-based inhibition effect was observed by Bäuml (1998) using a somewhat different paradigm.

There was a study by Jakab and Raaijmakers (2009) that tried to produce an analogue to this effect but failed to find it. However, in all of their experiments, the items were of medium category frequency. In the first two experiments, strength was defined in terms of the serial order in which items were presented in the initial study phase, and in the third experiment the strong items were presented twice, and the weak items were presented once. However, it is unclear whether serial position produces interference sufficient to recruit detectable differences in suppression, and why presenting items twice rather than once would produce meaningful differences to generate interference effects dramatic enough to be comparable with the substantial preexisting category strength differences a person may have. Moreover, there is no explanation of why differential inhibition effects for strong and weak items were produced by the 148 participants in the three Anderson et al. (1994) experiments other than an appeal to it as a "chance result". Thus, given these questionable methodological assumptions and the failure to find differential inhibition effects, these data were not considered here.

In terms of competitor set size interference, greater numbers of memory traces can produce greater levels of interference, which then need to be countered by suppression. While activation-only models can account for the greater interference, they do not account for the subsequent negative priming effect. The issue here is whether hydrogen would predict greater inhibition effects when there are more competitors. The answer is that it does not, but instead it predicts the opposite. To derive this prediction, we used the simulations from the negative priming paradigm described earlier and selected out those trials on which there were either one or two competitors (fan levels two and three) and involved the same target items with primes that had either two or three competitors. Thus, the number of associations was the same in the one- and two-distractor conditions. The results are summarized in Figure 12, which are graphed as the degree of retrieval suppression (experimental minus control) to make the size of the effects clearer. As can be seen, the inhibition effect was smaller when there were more competitors.

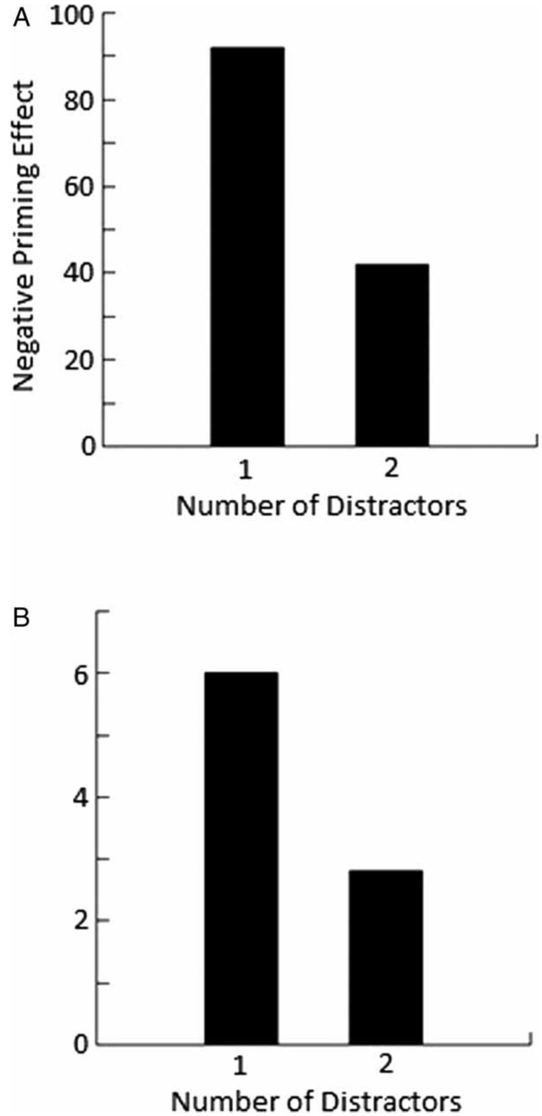

Figure 12. *Results from the negative priming experiments and simulation in terms of the number of distractors. A is the human data, and B is hydrogen with suppression operating.*

The simulation data were compared with the pattern of inhibition effects from similar studies (Radvansky, 1999b; Radvansky et al., 2005). These effects are also shown in Figure 12 using the averaged data across the studies. This analysis of the human data resulted in a large number of participants whose data needed to be excluded.

This happened because, in order to keep everything properly balanced, a person needed to have at least one correct response in all four trial types (experimental and control prime and target trials), and there were few trials per subject in each of these four trial types (*N* of 2 per condition, per participant). Note that some of these conditions are ones that are likely to produce larger errors, particularly in the experimental condition of the negative priming analysis. This small number of usable trials per participant is further driven by the fact that only a few of the targets in the negative priming paradigm had primes that had both one and two competitors. Without these constraints, the comparison is suspect and is probably not valid. Because we are dealing with an analysis that divides the data very finely resulting in a small number of trials and a number of excluded subjects, we collapsed the data across the three experiments, giving us 57 usable subjects, whose data are reported in Figure 12. As can be seen, both in the simulation and in the observed data, the size of the inhibition effect is smaller when there are more competitors (greater interference).

This result seems surprising in light of the data from the retrieval practice paradigm. Why is the size of the inhibition effect the opposite to that of the interference effect? The solution underscores the value of using approaches like hydrogen for more specific quantitative predictions. The reason for this pattern is because the amount of facilitation allocated to the target trace reduced with more traces in the retrieval set, and there is also a reduction in the amount of suppression applied to each of the competitors. So, while retrieval may take longer on the prime trial (producing the fan effect) for the memory to be retrieved, the absolute level of suppression for each competitor is reduced, causing retrieval time on the target trial to be less affected (the negative priming effect). This differential suppression is shown in Figure 13 in which the suppression levels of irrelevant memory traces are plotted over time for conditions in which there are one and two competitors (fan levels 2 and 3). As can be clearly seen, the single competitor is suppressed more than the dual-competitor trace.

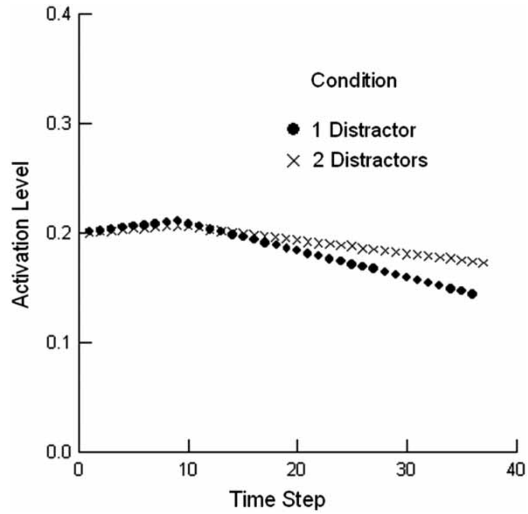

Figure 13. *Hydrogen simulation activation values for distractor traces from the fan effect/negative priming paradigm as a function of whether there are one or two distractors (fan levels 2 and 3). Note that because suppression is divided among the distractors, there is less suppression when there are two distractors than when there is only one distractor. Note that suppression starts operating after Time Step 10 in this simulation.*

So, based on this analysis, approaches such as hydrogen allow for the exploration and explanation of new issues. This analysis provides an account of the influence of the level of interference on the observed inhibition in the repeated practice paradigm, a prediction of the impact of levels of interference in a fan effect/negative priming paradigm, a confirmation of this prediction by an examination of the data, and an account for why these two patterns of data differ.

GENERAL DISCUSSION

The aim of the current project was to assess the utility of active suppression as a means of regulating interference during memory retrieval. This is in comparison to retrieval that only involves activation. In large measure, the claims on both sides have been made using verbal arguments, often focusing on the retrieval practice paradigm. There has been no explicit and direct attempt to back up a comparison of these ideas with the quantitative modelling that is found in other lines of memory research. The current study was a first step in that direction, albeit with a simplistic model. As a first step, some general claims about the operation of an active inhibitory mechanism can be addressed.

First, in all three paradigms tested—namely, the retrieval practice, part-set cueing, and fan effect/negative priming paradigms—we failed to find consistent strong evidence for a clear inhibition effect in the absence of suppression. That is, retrieval using activation only was influenced by the level of interference present (e.g., Anderson et al., 1994) when assessing whether an inhibition *effect* might emerge for memories with a raised activation level. Thus, it was more likely that they would be recalled later and rendered the other memory traces less likely to be recalled. While intuitively plausible, this pattern failed to consistently emerge. An inhibition effect was generated with an activation-only process in the part-set cueing paradigm, and then only when free recall was used. When recognition was used in this paradigm, or in the retrieval practice or fan effect/negative priming paradigms, the results obtained with an activation-only model were the opposite of what was actually observed. Thus, at this point, given the sort of model assessed here, the idea that these inhibition effects can be reliably produced with an activation-only based model is not supported.

In comparison, when suppression was operating in hydrogen, then inhibition effects clearly emerged. In these simulations, the same retrieval processes operated under both conditions, it is just that in one case suppression was turned off, and in the other it was turned on. Finally, using a model that includes active suppression allowed us to make predictions about the pattern of inhibition effects when different interference levels are generated by either the strength of the competitors or the number of competitors.

It should be noted that these results were obtained when (a) there was no parameter adjustment to better fit the model to the data, and (b) the model was applied to three different paradigms and (c) without any additional assumptions for each paradigm. With this approach, we found inhibitory effects using suppression. So, in terms of the basic retrieval processes used here, a model that includes suppression is preferable to one that does not.

One of the principles built into hydrogen is the idea that suppression is graded and limited. First, whether suppression is involved is a function of whether there is any interference. For hydrogen, retrieval suppression plays a particular role under particular circumstances. In structural suppression, suppression is part and parcel of the structure of the architecture and is always present. Also, with metacognitive suppression, as with directed forgetting, the process also operates in a way that suppression is always involved (e.g., if suppression is involved in directed forgetting, then it would be involved in all attempts to forget). In comparison, retrieval suppression only comes into play when there is a conflict to resolve—namely, interference. If there are no sources of interference, either because there are no other meaningfully related memory traces, or the information in the memory traces have been integrated into a common representation, then there is no reason to apply suppression, and so it is not operating.

Second, when interference is used, suppression is not applied uniformly under all conditions. Instead, it is limited to the amount of interference that is experienced. When interference levels are high, then more suppression is recruited to overcome that interference. However, when interference levels are low, then the amount of suppression operating is also low. This variability in the degree of suppression has been demonstrated in some of the work by M. C. Anderson and colleagues using the retrieval practice paradigm (e.g., Anderson et al., 1994; Anderson, Bjork, & Bjork, 2000; Anderson & Bell, 2001; Anderson & McCulloch, 1999; Anderson & Spellman, 1995; Levy & Anderson, 2002), and by Bäuml, Kissler, and Rak (2002) in the part-set cueing paradigm. Moreover, the amount of suppression to be distributed by a given memory trace is a function of how activated that trace is. This happens because the better match a trace is during the retrieval process, the greater the likelihood that it should be separated out and retrieved relative to the others.

Third, the amount of suppression that is available is divided up among the competitors, similar to what is done in explanations of the division of activation with the fan effect (e.g., Anderson, 1983). Thus, the effects of suppression are diffused by the needs of retrieval to deal with larger numbers of competitors, thereby further slowing retrieval and also reducing the amount that any given trace is suppressed.

One aspect of the hydrogen model that may seem objectionable to some readers is the absence of active suppression during free recall. We would like to make clear, again, that this does not mean that we argue for the absence of such a process during free recall. Instead, in an effort to try to keep things reasonably simple, we omitted it here, and we were able to simulate the basic free-recall processes that were used in these simulations. We acknowledge and are open to the possibility of an active suppression process during free recall as well. A more complex memory model other than hydrogen could incorporate such a process.

There are a number of directions that the hydrogen model can be taken. For instance, it could be applied to any of the different variations of the three research paradigms assessed here to either better understand those processes or to see when and how the model breaks down. One direction that we have taken is to make it part of a larger account of memory retrieval of event cognition in the event horizon model. This approach assesses how the segregation of information into events influences the retrieval information about those events later. Part of the theory incorporates principles of hydrogen to see how well it matches the data from various event cognition studies.

Specifically, the event horizon model uses the principles embedded in hydrogen to help explain why memory is better when event boundaries are present, as in some retroactive interference paradigms (Bilodeau & Schlossberg, 1951; Greenspoon & Ranyard, 1957; Jensen, Dibble, & Anderson, 1971; Nagge, 1935; Smith, Glenberg, & Bjork, 1978; Strand, 1970). Also, when there are multiple competing traces (event models in this case) in a recognition paradigm, the presence of event boundaries can make memory worse, as when people make recognition decisions about information that was part of multiple events as was either explicitly learned (Bower & Rinck, 2001; Radvansky, 1998, 1999b, 2005, 2009; Radvansky & Copeland 2006b; Radvansky et al., 1993; Radvansky & Zacks, 1991) or when information about an event is transferred from one event to another (Radvansky & Copeland, 2006b; Radvansky, Tamplin, & Krawietz, 2010; Swallow, Zacks, & Abrams, 2009).

So, in conclusion, we have demonstrated the viability of suppression as a mechanism of memory retrieval across a number of paradigms, and the comparison with an activation-only account was done by not altering the simulations other than to allow suppression to operate or not. This approach can be further tested in the future to determine its capabilities and limitations. It can also be applied to future investigations of the impact of retrieval interference and suppression on other cognitive phenomena.

Original manuscript received 19 May 2011
Accepted revision received 22 August 2012
First published online 23 November 2012

REFERENCES

Anderson, J. R. (1974). Retrieval of propositional information from longterm memory. *Cognitive Psychology*, *6*, 451–474.

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Anderson, J. R., & Reder, L. M. (1999). The fan effect: New results and new theories. *Journal of Experimental Psychology: General*, *128*, 186–197.

Anderson, M. A., Bjork, R. A., & Bjork, E. L. (1994). Remembering can cause forgetting: Retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 1063–1087.

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

Anderson, M. C., & Bell, T. (2001). Forgetting our facts: The role of inhibitory processes in the loss of propositional knowledge. *Journal of Experimental Psychology: General*, *130,* 544–570.

Anderson, M. C., & McCulloch, K. C. (1999). Integration as a general boundary condition on retrieval-induced forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25,* 608–629.

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

Aslan, A., & Bäuml, K.-H. (2007). Part-list cuing with and without item-specific probes: The role of encoding. *Psychonomic Bulletin & Review*, *14,* 489–494.

Aslan, A., & Bäuml, K.-H. T. (2010). Retrieval-induced forgetting in young children. *Psychonomic Bulletin & Review*, *17,* 704–709.

Aslan, A., & Bäuml, K.-H. (2011). Individual differences in working memory capacity predict retrieval-induced forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37,* 264–269.

Aslan, A., Bäuml, K.-H., & Grundgeiger, T. (2007). The role of inhibitory processes in part-list cuing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33,* 335–341.

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

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

Barnier, A. J., Hung, L., & Conway, M. A. (2004). Retrieval-induced forgetting of emotional and unemotional autobiographical memories. *Cognition and Emotion*, *18,* 457–477.

Basden, D. R., & Basden, B. H. (1995). Some tests of the strategy disruption interpretation of part-list cuing inhibition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21,* 1656–1669.

Basden, D. R., Basden, B. H., & Galloway, B. C. (1977). Inhibition with part-list cuing: Some tests of the item strength hypothesis. *Journal of Experimental Psychology: Human Learning and Memory*, *3,* 100–108.

Bäuml, K.-H. (1998). Strong items get suppressed, weak items do not: The role of item strength in output interference. *Psychonomic Bulletin & Review*, *5,* 459–463.

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

Bäuml, K.-H., & Aslan, A. (2004). Part-list cuing as instructed retrieval inhibition. *Memory & Cognition*, *32,* 610–617.

Bäuml, K.-H., & Aslan, A. (2006). Part-list cuing can be transient and lasting: The role of encoding. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32,* 33–43.

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

Bäuml, K.-H., Kissler, J., & Rak, A. (2002). Part-list cuing in amnesic patients: Evidence for a retrieval deficit. *Memory & Cognition*, *30,* 862–870.

Bäuml, K.-H., & Kuhbander, C. (2003). Retrieval-induced forgetting and part-list cuing in associatively structured lists. *Memory & Cognition*, *31,* 1188–1197.

Bäuml, K.-H., & Kuhbander, C. (2007). Remembering can cause forgetting—but not in negative moods. *Psychological Science*, *18,* 111–115.

Bäuml, K.-H., Zillner, M., & Vilimek, R. (2005). When remembering causes forgetting: Retrieval-induced forgetting as recovery failure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31,* 1221–1234.

Bilodeau, I. M., & Schlossberg, H. (1951). Similarity in stimulus conditions and retroactive inhibition. *Journal of Experimental Psychology*, *41,* 199–204.

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

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

Bower, G. H., & Rinck, M. (2001). Selecting one among many referents in spatial situation models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27,* 81–98.

Brown, A. S., & Hall, L. A. (1979). Part-list cueing inhibition in semantic memory structures. *American Journal of Psychology*, *92,* 351–362.

Butler, K. M., Williams, C. C., Zacks, R. T., & Maki, R. H. (2001). A limit on retrieval-induced forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27,* 1314–1319.

Camp, G., Pecher, D., & Schmidt, H. G.(2005). Retrieval-induced forgetting in implicit memory

tests: The role of test awareness. *Psychonomic Bulletin & Review*, *12*, 490–494.
Camp, G., Pecher, D., & Schmidt, H. G. (2007). No retrieval-induced forgetting using item-specific independent cues: Evidence against a general inhibitory account. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 950–958.
Carlson, N. R. (2004). *Foundations of physiological psychology*. Boston, MA: Allyn & Bacon.
Christensen, B. K., Girard, T. A., Benjamin, A. S., & Vidailhet, P. (2006). Evidence for impaired mnemonic strategy use among patients with schizophrenia using the part-list cuing paradigm. *Schizophrenia Research*, *85,* 1–11.
Cokely, E. T., Kelley, C. M., & Gilchrist, A. L. (2006). Sources of individual differences in working memory: Contributions of strategy to capacity. *Psychonomic Bulletin & Review*, *13,* 991–997.
Crescentini, C., Shallice, T., Del Missier, F., & Macaluso, E. (2010). Neural correlates of episodic retrieval: An fMRI study of the part-list cueing effect. *NeuroImage*, *50,* 678–692.
Dewhurst, S. A., Bould, E., Knott, L. M., & Thorley, C. (2008). The roles of encoding and retrieval processes in associative and categorical memory illusions. *Journal of Memory and Language*, *60,* 154–164.
Dodd, M. D., Castel, A. D., & Roberts, K. E. (2006). A strategy disruption component to retrieval-induced forgetting. *Memory & Cognition*, *34,* 102–111.
Greenspoon, J., & Ranyard, R. (1957). Stimulus conditions and retroactive inhibition. *Journal of Experimental Psychology*, *53,* 55–59.
Groome, D., & Grant, N. (2005). Retrieval-induced forgetting is inversely related to everyday cognitive failures. *British Journal of Psychology*, *96,* 313–319.
Hanslmayr, S., Staudigl, T., Aslan, A., & Bäuml, K.-H. (2010). Theta oscillations predict the detrimental effects of memory retrieval. *Cognitive, Affective, & Behavioral Neuroscience*, *10,* 329–338.
Healey, M. K., Campbell, K. L., Hasher, L., & Ossher, L. (2009). Direct evidence for the role of inhibition in resolving interference in memory. *Psychological Science*, *21,* 1464–1470.
Hicks, J. L., & Starns, J. J. (2004). Retrieval-induced forgetting occurs in tests of item recognition. *Psychonomic Bulletin & Review*, *11*, 125–130.
Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, *95,* 528–551.
Jakab, E., & Raaijmakers, J. G. W. (2009). The role of item strength in retrieval-induced forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35,* 607–617.
Jensen, L., Dibble, J., & Anderson, D. C. (1971). Effects of a contextual change upon retroactive inhibition. *Psychological Reports*, *29,* 39–46.
Johansson, M., Aslan, A., Bäuml, K.-H., Gäbel, A., & Mecklinger, A. (2007). When remembering causes forgetting: Electrophysiological correlates of retrieval-induced forgetting. *Cerebral Cortex*, *17,* 1335–1341.
Koessler, S., Engler, H., Riether, C., & Kissler, J. (2009). No retrieval-induced forgetting under stress. *Psychological Science*, *20,* 1356–1363.
Kuhl, B. A., Dudukovic, N. M., Kahn, I., & Wagner, A. D. (2007). Decreased demands on cognitive control reveal the neural processing benefits of forgetting. *Nature Neuroscience*, *10,* 908–914.
Levy, B. J., & Anderson, M. C. (2002). Inhibitory processes and the control of memory retrieval. *Trends in Cognitive Science*, *6,* 299–305.
Luchuga, M. T., Moreno, V., Pelegrina, S., Gómez-Ariza, C. J., & Bajo, M. T. (2006). Age differences in memory control: Evidence from updating and retrieval-practice tasks. *Acta Psychologica*, *123,* 279–298.
MacLeod, C. M., Dodd, M. D., Sheard, E. D., Wilson, D. E., & Bibi, U. (2003). In opposition to inhibition. *Psychology of Learning and Motivation*, *43,* 163–214.
MacLeod, C. M., & Macrae, C. N. (2001). Gone but not forgotten: The transient nature of retrieval-induced forgetting. *Psychological Science*, *12,* 148–152.
Marsh, E. J., Dolan, P. O., Balota, D. A., & Roediger, H. L. (2004). Part-set cuing effects in younger and older adults. *Psychology and Aging*, *19,* 134–144.
Mueller, C. W., & Watkins, M. J. (1977). Inhibition from part-set cuing: A cue-overload interpretation. *Journal of Verbal Learning and Verbal Behavior*, *16,* 699–709.
Nagge, J. W. (1935). An experimental test of the theory of associative interference. *Journal of Experimental Psychology*, *18,* 663–682.
Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, *106,* 226–254.
Nestor, P. G., Piech, R., Allen, C., Niznikiewicz, M., Shenton, M., & McCarley, R. W. (2005). Retrieval-induced forgetting in schizophrenia. *Schizophrenia Research*, *75,* 199–209.

Nievas, F., & Mari-Beffa, P. (2002). Negative priming from the non-selected meaning of the homograph. *British Journal of Psychology*, *93,* 47–66.

Norman, K. A., Newman, E. L., & Detre, G. (2007). A neural network model of retrieval-induced forgetting. *Psychological Review*, *114,* 887–953.

Oswald, K. M., Serra, M., & Krishna, A. (2005). Part-list cuing in speeded recognition and free recall. *Memory & Cognition*, *34,* 518–526.

Park, H., & Madigan, S. (1993). Negative effects of part-set cues: Are they reversible? *Bulletin of the Psychonomic Society*, *31,* 311–313.

Perfect, T. J., Moulin, C. J. A., Conway, M. A., & Perry, E. (2002). Assessing the inhibitory account of retrieval-induced forgetting with implicit-memory tests. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*, 1111–1119.

Perfect, T. J., Stark, L.-J., Tree, J. J., Moulin, C. J. A., Ahmed, L., & Hutter, R. (2004). Transfer appropriate forgetting: The cue-dependent nature of retrieval-induced forgetting. *Journal of Memory and Language*, *51,* 399–417.

Raaijmakers, J. G. W., & Phaf, R. M. (1999). Part-list cuing revisited: Testing the sampling-bias hypothesis. In C. Izawa (Ed.), *On human memory: Evolution, progress, and reflections on the 30th anniversary of the Atkinson–Shiffrin model* (pp. 87–104). Mahwah, NJ: Lawrence Erlbaum Associates.

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

Radvansky, G. A. (1998). The organization of information retrieved from situation models. *Psychonomic Bulletin & Review*, *5,* 283–289.

Radvansky, G. A. (1999a). The fan effect: A tale of two theories. *Journal of Experimental Psychology: General*, *128,* 198–206.

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

Radvansky, G. A. (2005). Situation models, propositions, and the fan effect. *Psychonomic Bulletin & Review*, *12,* 478–483.

Radvansky, G. A. (2009). Spatial directions and situation model organization. *Memory & Cognition*, *37*, 796–806.

Radvansky, G. A., & Copeland, D. E. (2006a). Memory retrieval and interference: Working memory issues. *Journal of Memory and Language*, *55,* 33–46.

Radvansky, G. A., & Copeland, D. E. (2006b). Walking through doorways causes forgetting. *Memory & Cognition*, *34,* 1150–1156.

Radvansky, G. A., Spieler, D. H., & Zacks, R. T. (1993). Mental model organization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19,* 95–114.

Radvansky, G. A., Tamplin, A. K., & Krawietz, S. A. (2010). Walking through doorways causes forgetting: Environmental integration. *Psychonomic Bulletin & Review*, *17,* 900–904.

Radvansky, G. A., & Zacks, R. T. (1988). [Fan effect]. Unpublished raw data.

Radvansky, G. A., & Zacks, R. T. (1991). Mental models and the fan effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17,* 940–953.

Radvansky, G. A., Zacks, R. T., & Hasher, L. (1996). Fact retrieval in younger and older adults: The role of mental models. *Psychology and Aging*, *11,* 258–271.

Radvansky, G. A., Zacks, R. T., & Hasher, L. (2005). Age and inhibition: The retrieval of situation models. *Journal of Gerontology: Psychological Sciences*, *60B*, P276–P278.

Reysen, M. B., & Nairne, J. S. (2002). Part-set cuing of false memories. *Psychonomic Bulletin & Review*, *9,* 389–393.

Rhodes, M. G., & Castel, A. D. (2008). Metacognition and part-set cuing: Can interference be predicted at retrieval? *Memory & Cognition*, *36,* 1429–1438.

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

Romàn, P., Soriano, M. F., Gómez-Ariza, C. J., & Bajo, M. T. (2009). Retrieval-induced forgetting and executive control. *Psychological Science*, *20,* 1053–1058.

Rundus, D. (1973). Negative effects of using list items as recall cues. *Journal of Verbal Learning and Verbal Behavior*, *12,* 43–50.

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

Slamecka, N. J. (1968). An examination of trace storage in free recall. *Journal of Experimental Psychology*, *4,* 504–513.

Slamecka, N. J. (1972). The question of associative growth in the learning of categorized material. *Journal of Verbal Learning and Verbal Behavior*, *11,* 324–332.

Sloman, S. A., Bower, G. H., & Rohrer, D. (1991). Congruency effects in part-list cuing inhibition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 974–982.

Smith, S. M., Glenberg, A., & Bjork, R. A. (1978). Environmental context and human memory. *Memory & Cognition*, *6,* 342–353.

Sohn, M.-H., Anderson, J. R., Reder, L. M., & Goode, A. (2004). Differential fan effect and attentional focus. *Psychonomic Bulletin & Review*, *11,* 729–734.

Soriano, M. F., Jiménez, J. F., Romàn, P., & Bajo, M. T. (2009). Inhibitory processes in memory are impaired in schizophrenia: Evidence from retrieval induced forgetting. *British Journal of Psychology*, *100,* 661–673.

Spitzer, B., & Bäuml, K.-H. (2007). Retrieval-induced forgetting in item recognition: Evidence for a reduction in general memory strength. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33,* 863–875.

Spitzer, B., Hanslmayr, S., Opitz, B., Mecklinger, A., & Bäuml, K.-H. (2008). Oscillatory correlates of retrieval-induced forgetting in recognition memory. *Journal of Cognitive Neuroscience*, *21,* 976–990.

Staudigl, T., Hanslmayr, S., & Bäuml, K.-H. (2010). Theta oscillations reflect the dynamics of interference in episodic memory retrieval. *Journal of Neuroscience*, *30,* 11356–11362.

Strand, B. Z. (1970). Change of context and retroactive inhibition. *Journal of Verbal Learning and Verbal Behavior*, *9,* 202–206.

Swallow, K. M., Zacks, J. M., & Abrams, R. A. (2009). Event boundaries in perception affect memory encoding and updating. *Journal of Experimental Psychology: General*, *138*, 236–257.

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

Todres, A. K., & Watkins, M. J. (1981). A part-set cuing effect in recognition memory. *Journal of Experimental Psychology: Human Learning and Memory*, *7*, 91–99.

Williams, C. C., & Zacks, R. T. (2001). Is retrieval-induced forgetting an inhibitory process? *The American Journal of Psychology*, *114,* 329–354.

Wimber, M., Bäuml, K.-H., Bergström, Z., Markopoulos, G., Heinze, H.-J., & Richardson-Klavehn, A. (2008). Neural markers of inhibition in human memory retrieval. *The Journal of Neuroscience*, *10*, 13419–13427.

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

Zellner, M., & Bäuml, K.-H. (2005). Intact retrieval inhibition in children's episodic recall. *Memory & Cognition*, *33,* 396–404.

APPENDIX A

Formulae for facilitation and other processes

During retrieval there is a facilitation influence on each trace, F_i. The facilitation to memory traces prior to the introduction of suppression:

$$F_i = \sum_{1}^{e} \frac{(1/e)}{N_{Match}}$$

where e is the number of components in the probe, and N_{Match} is the number of traces that match for a given component. The activation level of a memory trace, i, is denoted as A_i. The activation level of a trace after each time step is:

$$A_i = A_i + (F_i \times R)$$

Here R is the rate at which the simulation is progressing. For all of the simulations reported here, R was set to .005.

After a period of time, T, the level of interference is determined. This is the difference in activation levels of the traces in the retrieval set. A_i is the activation value of a given trace, and A_j is that for all of the other traces in the set; *NumComp* corresponds to the number of trace comparisons in a retrieval set and would equal $N \times (N-1)/2$. Interference is denoted as K:

$$K = 1 - \left(\frac{\sum_{i=1}^{N-1} \sum_{j=i+1}^{N} abs(A_i - A_j)}{NumComp} \right)$$

The level of suppression is a function of the level of interference. Thus, the inhibition level, I, where F_j is the facilitation level of every trace in the retrieval set, except for the current trace i, is

$$I_i = \sum_{j=1}^{N} \left(\frac{F_j}{(N-1)} \right)$$

Inhibition is divided in a manner similar to the division of facilitation based on the number of traces and the number of probe components involved, with the level of inhibition for a trace being the sum of the inhibitory inputs from its competitors. This can be described as

$$I_i = K \times \sum_{j=1}^{N} \left(\frac{F_j}{N-1} \right)$$

After the inhibition levels are set, both the facilitation and inhibition continue to contribute to a memory trace's activation level. Thus, the formula becomes

$$A_i = A_i + [(F_i - I_i) \times R]$$

This retrieval process continues, until a trace reaches the criteria level, C, relative to all of the traces, and has at least exceeded the criteria level relative to the baseline level. For the simulations reported here it is set at .1.

Those traces that have either been recently activated or inhibited decay back to this baseline via a negatively accelerating power function. The decay rate is dictated by the exponent in the equation, which captures how far the current activation level is from the baseline. Moreover, this is modified by a decay constant, D, which is a number between 0 and 1, with values such as .1 or .01 providing adequate speed of decay. The formula for applying the decay function is:

$$A_i = A_i - (D \times R)^{(1 = |A_i - Baseline|)}$$

Retrieval and negative responses

Essentially, after this adjusted period of time, if the top activation level in the retrieval set is still not sufficiently beyond the criterion at that point in time, then retrieval stops.

$$A_{\max} < Baseline + (1 - \left(\frac{Adjustment}{C} \right) \times C^{[C/(t \times F_{\max} \times R)]}$$

In this formula, adjustment is equal to $(.1 \times C)$ if there is only a single trace in the retrieval set, but is equal to the largest difference between activation levels, unless that difference is greater than .95, in which case adjustment is equal to .95.

APPENDIX B

Human data for the retrieval practice paradigm

Table B1. *Final recall test*

Study	*Exp*	N	*RP+*	*RP–*	*NRP*	*Notes*
Anderson et al. (1994)	1	36	.73	.38	.49	
	2	48	.85	.69	.73	
	3	64	.72	.49	.52	
Anderson & McCulloch (1999)	1	144	.73	.37	.45	
	2	96	.76	.43	.50	
	3	128	.87	.65	.70	
Anderson & Spellman (1995)	1	48	.69	.24	.38	
	2	54	.72	.42	.52	
	3B	48	.72	.35	.47	
	4	72	.70	.43	.57	
Aslan & Bäuml (2010)	1	48	.81	.53	.60	Adults only
Aslan, Bäuml, & Pastotter (2007)	1	24	.80	.56	.62	Younger adults only
Bäuml & Hartinger (2002)	1	54	.96	.85	.90	
Bäuml & Kuhbander (2007)	1	27	.79	.49	.57	Neutral mood only
Bäuml, Zillner, & Vilimek (2005)	1	24	.89	.59	.69	
	2	24	.76	.48	.59	
Butler, Williams, Zacks, & Maki (2001)	1A	36	.73	.42	.47	
	1B	36	.56	.45	.45	
	1C	36	.66	.50	.51	
	1D	36	.90	.79	.80	
	1E	36	.90	.85	.82	
Camp, Pecher, & Schmidt (2007)	1	36	.81	.55	.63	Estimated
	2	30	.62	.49	.40	Estimated
	3	30	.57	.38	.32	Estimated
Dodd, Castel, & Roberts (2006)	1	24	.71	.23	.33	Random order condition
	2	24	.71	.28	.33	Random order condition
	3	20	.83	.54	.65	Random order condition
Groome & Grant (2005)	1	40	.86	.42	.55	
Hanslmayr et al. (2010)	1	24	.85	.58	.64	Competitive condition
Jakab & Raaijmakers (2009)	1	51	.71	.37	.41	
	2	51	.73	.37	.45	
	3	41	.59	.30	.41	
Koessler, Engler, Riether, & Kissler (2009)	1	24	.70	.44	.51	Controls only
Kuhl et al. (2007)	1	24	.48	.30	.35	
Lechuga, Moreno, Pelegrina, Gómez-Ariza, & Bajo (2006)	2	30	.55	.28	.40	Younger adults only
Macleod & Macrae (2001)	1	64	.72	.16	.30	Only immediate not 24-hour delay
	2	64	.68	.18	.32	
Nestor et al. (2005)	2	16	.74	.39	.46	Controls only
Perfect et al. (2004)	1	20	.96	.53	.71	
	2	18	.99	.44	.61	
Saunders, Fernandes, & Kosnes (2009)	1	21	.90	.24	.41	
Soriano, Jiménez, Romàn, & Bajo (2009)	1	18	.58	.18	.28	Controls only
Williams & Zacks (2001)	1	48	.70	.38	.44	
	2	48	.66	.39	.46	
	3	90	.72	.42	.48	
Wimber et al. (2008)	1	23	.70	.53	.53	
Weighted mean			*.75*	*.44*	*.52*	

Note: Exp = experiment. RP + =items practised in the study phase; RP– = items not practised during retrieval practice, but from the same categories; NRP = items from categories that were only studied and not practised at all.

Table B2. *Final recognition test*

Study	*Exp*	N	*RP+*	*RP–*	*NRP*	*Notes*
Aslan & Bäuml (2010)	1	48	.99	.64	.76	Adults only
Aslan & Bäuml (2011)	1	168	.90	.64	.69	
Hicks & Starns (2004)	1	40	.91	.62	.71	
Romàn, Soriano, Gómez-Ariza, & Bajo (2009)	1	32	.74	.49	.62	Controls only
Soriano et al. (2009)	2	22	.90	.56	.65	Controls only
Spitzer & Bäuml (2007)	1	48	.85	.60	.67	
Spitzer, Hanslmayr, Opitz, Mecklinger, & Bäuml (2008)	1	30	.96	.77	.84	
Weighted mean			*.90*	*.63*	*.70*	

Note: Exp = experiment. RP + =items practised in the study phase; RP– = items not practised during retrieval practice, but from the same categories; NRP = items from categories that were only studied and not practised at all.

APPENDIX C

Human data from various part-set cueing studies

Table C1. *Final recall test*

Study	*Exp*	N	*Control*	*Experimental*	*Notes*
Bäuml & Aslan (2006)	1	72	.80	.69	Collapsed across high and low associativity
Bäuml & Kuhbander (2003)	1	56	.64	.53	
	2	66	.60	.45	
Brown & Hall (1979)	1	75	.71	.63	
Christensen, Girard, Benjamin, & Vidailhet (2006)	1	30	.56	.52	Scrambled
	1	30	.72	.64	Blocked
Cokely, Kelley, & Gilchrist (2006)	1	40	.43	.33	
	2	64	.37	.31	
Dewhurst, Bould, Knott, & Thorley (2008)	4	56	.61	.56	
Marsh, Dolan, Balota, & Roediger (2004)	1	27	.74	.62	Younger adults only
	2	27	.60	.52	
	3	27	.61	.61	
Mueller & Watkins (1977)	1	32	.58	.48	
	3	21	.62	.53	
	4	51	.61	.57	
Oswald et al. (2005)	1	144	.77	.69	
	2	72	.55	.47	
Park & Madigan (1993)	1	44	.43	.30	Collapsed across A and B—first test only
	2	60	.40	.36	First test only
Raaijmakers & Phaf (1999)	1	25	.46	.48	No cues and random conditions
Reysen & Nairne (2002)	1	62	.51	.45	
	2	84	.50	.42	
Rhodes & Castel (2008)	2	60	.39	.30	
	3	48	.46	.33	
Roediger (1973)	1	48	.66	.62	
Rundus (1973)	1	11	.38	.40	
	2	13	.35	.31	
Slamecka (1972)	2	30	.48	.71	
Sloman, Bower, & Rohrer (1991)	1	28	.21	.20	
Zellner & Bäuml (2005)	2	24	.88	.77	Younger adults only
Weighted mean			*.57*	*.51*	

Note: Exp = experiment. RP + =items practised in the study phase; RP– = items not practised during retrieval practice, but from the same categories; NRP = items from categories that were only studied and not practised at all.

Table C2. *Final recognition test*

Study	*Exp*	N	*Control*	*Experimental*	*Notes*
Oswald et al. (2005)	1	144	.90	.88	
Slamecka (1972)	1	18	.83	.84	Collapsed across trial
	2	24	.66	.68	Collapsed across trial
Todres & Watkins (1981)	1	54	.82	.81	
	3	56	.79	.79	
	4	56	.75	.72	
Weighted mean			*.83*	*.81*	

Note: Exp = experiment. RP + =items practised in the study phase; RP− = items not practised during retrieval practice, but from the same categories; NRP = items from categories that were only studied and not practised at all.

APPENDIX D

Data from fan effect experiments

Table D1. *Human data from fan effect experiments*

			Single location			*Multiple locations*			
Study	*Exp*	N	*1*	*2*	*3*	*1*	*2*	*3*	*Notes*
Radvansky & Zacks (1988)		32	1,670	1,613	1,533	1,670	1,766	1,963	No fan level 4
Radvansky & Zacks (1991)	1	32	1,536	1,623	1,608	1,536	1,787	1,755	No fan level 4
	2	72	1,513	1,484	1,534	1,513	1,657	1,775	No fan level 4
	3	72	1,443	1,414	1,388	1,443	1,509	1,624	No fan level 4. negative priming
Radvansky & Zacks (1991)		32	1,466	1,379	1,365	1,421	1,592	1,772	
Radvansky et al. (1993)	1	48	1,424	1,502	1,477	1,393	1,550	1,706	
	2	48	1,389	1,426	1,386	1,433	1,462	1,574	
Radvansky et al. (1996)	1	28	1,558	1,524	1,552	1,541	1,796	1,780	Younger adults
	2	32	1,346	1,390	1,344	1,387	1,531	1,687	Younger adults
	3	32	1,578	1,632	1,573	1,550	1,627	1,594	Younger adults
Radvansky (1998)		48	1,565	1,581	1,530	1,531	1,671	1,748	
Radvansky (1999b)	1	71	1,665	1,580	1,581	1,645	1,863	1,893	No NP targets
Radvansky (1999b)	2	48	1,387	1,347	1,502	1,316	1,469	1,659	No NP targets
Radvansky (1999b)	3	48	1,327	1,357	1,433	1,367	1,464	1,489	
Radvansky (2005)		66	1,404	1,354	1,382	1,401	1,570	1,592	Sentences
Radvansky & Copeland (2006a)		150	1,452	1,476	1,468	1,459	1,662	1,778	
Radvansky & Copeland (2006b)	1	36	1,263	1,362	1,265	1,297	1,520	1,483	
Weighted mean			*1,457*	*1,466*	*1,462*	*1,454*	*1,610*	*1,689*	

Note: Exp = experiment. NP = negative priming. Values are in milliseconds.

Table D2. *Data from negative priming target trials of fan effect experiments*

Study	*Exp*	N	*Control*	*Experimental*	*Notes*
Radvansky (1999b)	1	71	1,676	1,736	
	2	48	1,470	1,595	
Radvansky et al. (2005)		32	1,755	1,875	Younger adults
Radvansky & Copeland (2006a)		150	1,550	1,589	
Weighted mean			*1,591*	*1,655*	

Note: Exp = experiment. Values are in milliseconds.

APPENDIX E

Human data from retrieval practice experiments in which items varied in taxonomic category strength

Table E1. *Human data from retrieval practice experiments*

Study	*Exp*	N	*Strong effect*	*Weak effect*
Anderson et al. (1994)	1	36	.16	.06
	2	48	.08	.00
	3	64	.09	−.06
Williams & Zacks (2001)	1	48	.07	.06
	2	48	.06	.08
Weighted mean			*.09*	*.02*

Note: Exp = experiment.