

On the Status of Inhibitory Mechanisms in Cognition: Memory Retrieval as a Model Case

Michael C. Anderson
University of California, Los Angeles

Barbara A. Spellman
University of Texas at Austin

Theories of cognition frequently assume the existence of inhibitory mechanisms that deactivate mental representations. Justifying this assumption is difficult because cognitive effects thought to reflect inhibition can often be explained without recourse to inhibitory processes. This article addresses the uncertain status of cognitive inhibitory mechanisms, focusing on their function in memory retrieval. On the basis of a novel form of forgetting reported herein, it is shown that classical associative theories of interference are insufficient as accounts of forgetting and that inhibitory processes must be at work. It is argued that inhibitory processes are used to resolve computational problems of selection common to memory retrieval and selective attention and that retrieval is best regarded as conceptually focused selective attention.

Advances in both brain and computational sciences have fueled an explosion of theories positing inhibitory mechanisms in a variety of cognitive functions. Theories of intact cognition emphasize the importance of these activation-reducing mechanisms to basic capacities, including visual selective attention (Duncan & Humphreys, 1989; Keele & Neill, 1978; Neill & Westberry, 1987; Neumann & DeSchepper, 1992; Tipper, 1985; Tipper, MacQueen, & Brehaut, 1988), language comprehension and production (Gernsbacher, 1990, 1991; Gernsbacher & Faust, 1991; Simpson & Kang, 1994), and retrieval from episodic and semantic memory (Anderson, Bjork, & Bjork, 1994; Bjork, 1989; Blaxton & Neely, 1983; Carr & Dagenbach, 1990), as well as to some higher order functions such as analogical reasoning (Holyoak & Thagard, 1989; Spellman & Holyoak, 1993)

and executive control (Logan & Cowan, 1984). Cognitive changes associated with normal aging (Hasher & Zacks, 1988; McDowd, Oseas-Kreger, & Fillion, 1994) and with several clinical syndromes, including schizophrenia (Beech, Powell, McWilliams, & Claridge, 1989; Cohen & Servan-Schreiber, 1992) and frontal-lobe dysfunction (Fuster, 1989; Luria, 1966; Mishkin, 1964), have been attributed to deficits in inhibitory function. The ability to inhibit extraneous information has even been proposed as an important dimension of general intelligence (Dempster, 1991; see Dagenbach & Carr, 1994a, and Dempster & Brainerd, 1994, for collections of reviews on inhibitory processes). The existence of such inhibitory mechanisms in the functional architecture of cognition seems both plausible and necessary: plausible because the substrate on which that architecture operates—the brain—uses both excitatory and inhibitory processes to perform neural computation, and necessary because computational analyses show that inhibitory mechanisms are critical for maintaining stability in neuronal networks (e.g., Easton & Gordon, 1984).

Michael C. Anderson, Department of Psychology, University of California, Los Angeles; Barbara A. Spellman, Department of Psychology, University of Texas at Austin.

This work was done while the authors were graduate students together at the University of California, Los Angeles; it constitutes one chapter in Michael C. Anderson's dissertation, which appears on University Microfilms. The research reported herein was supported in part by Grant 4-564040-RB-19900 from the Committee on Research, University of California, Los Angeles, and by National Science Foundation Grant SBR-9310614.

We would like to thank a legion of ex-undergraduates for donating their time and names to our experiments. In particular, we would like to thank Anthony Wagner, Shirley Yu, Steve Machado, Lisa Kaufman, Barnett Hartston, and Fernando Reynega. We would also like to thank the members of CogFog and Mini-Fog for sitting through endless discussions of these issues. We thank Keith Holyoak and William Trammel Neill for helpful comments on an earlier draft and Randy Gallistel, Susan Bookheimer, and Robert Bjork for comments on the final version.

Correspondence concerning this article should be addressed to Michael C. Anderson, who is now at Department of Psychology, University of California, Berkeley, 3210 Tolman Hall, #1650 Berkeley, California, 94720. After September 1995, the address will be Department of Psychology, College of Arts and Sciences, University of Oregon, Eugene, Oregon 97403-1227.

Whereas many psychologists agree with the general theoretical plausibility of cognitive inhibition, consensus about the role of inhibition in any particular empirical effect has remained elusive. Consider, for example, the ambiguities in interpreting increased response latency as evidence for cognitive inhibition. Increases in the latency to respond to a stimulus (e.g., a picture or a word) suggest inhibition if one assumes that latencies directly measure the activation state of a representation. If, however, one recognizes that other factors, such as response conflict or subject decision strategies, might increase latencies when there is more activation, then increased response latencies can result from increased activation rather than inhibition. For instance, suppose that participants are slower to decide whether the probe word ACE is semantically congruent with the prime sentence "He dug with the spade" than with the sentence "He dug with the shovel." Slowed reaction time might result from inhibition of the contextually inappropriate sense of the word SPADE (i.e., playing card) during that word's disambiguation in the previously read sentence. Alternatively, slowed reaction

time to ACE might result from the contextually inappropriate sense of SPADE remaining highly active, causing decision conflict that increases latencies (Gernsbacher & Faust, 1991). Similar ambiguities beset evidence for inhibition in other domains, such as selective attention (Park & Kanwisher, 1994) and memory retrieval (Blaxton & Neely, 1983; Brown, 1981), leading to reasonable skepticism regarding the status of cognitive inhibition. When empirical ambiguities render it difficult to determine the appropriate account, postulation of inhibitory processes is left dependent solely on parsimony or theoretical taste. What is needed are methods that establish inhibitory mechanisms as the source of the "inhibitory" effects.

This article addresses the ambiguous status of cognitive inhibition in the domain of retrieval from long-term memory. Memory retrieval offers an excellent domain in which to test inhibitory theories for several reasons. First, the alternative noninhibitory theories of memory retrieval are important in themselves, are clearly articulated, and are testable. Second, a number of authors have argued lucidly for the function that inhibition might serve in memory retrieval (Blaxton & Neely, 1983; Carr & Dagenbach, 1990; Dagenbach & Carr, 1994b; Keele & Neill, 1978; Neill & Westberry, 1987). Finally, the role that inhibition is thought to play in memory retrieval is similar to the function that it is thought to serve in other domains such as language processing and selective attention. Indeed, several authors have noted the strong computational similarities shared by these domains (Anderson & Bjork, 1994; Neill, Valdes, & Terry, in press; Neill & Westberry, 1987; Neumann & DeSchepper, 1992; Posner, 1980, 1987). Thus, to the extent that claims of cognitive inhibition may be evaluated in the domain of memory retrieval, one may inform theories for other cognitive functions in which similar functional circumstances apply.

We begin our inquiry into inhibition by illustrating the theoretical and empirical reasons why the status of inhibitory processes in memory retrieval has been a source of debate. Next, we describe a unique prediction of inhibitory models that allows them to be distinguished from noninhibitory alternatives. On the basis of that prediction, we then develop the general method used in all of our attempts to resolve the status of cognitive inhibition.

INHIBITORY PROCESSES IN MEMORY RETRIEVAL: EXAMPLES AND DEBATE

Although the argument for inhibitory mechanisms in memory retrieval seems plausible, disagreement remains over the existence of such processes. In this section, we describe the theoretical case for the role of inhibitory mechanisms in memory retrieval. After providing examples of phenomena that appear to support such mechanisms, we illustrate the basis of the theoretical disagreement by showing how the same phenomena may be explained by several alternative noninhibitory theories.

Arguments for Inhibition

The argument for inhibitory processes in memory retrieval is plausible on both functional and empirical grounds. Consider first the function that inhibition might serve in the task of retrieving target items from long-term memory. Retrieval often

requires access of a specific fact or event given some fairly general retrieval cues to start. These general retrieval cues tend to activate a broad set of memory items before the target can be isolated. For example, presentation of the category label FRUIT would initially activate many exemplars of the Fruit category, even if only the item ORANGE is sought. When activation diffuses broadly in this manner, some mechanism that narrows the focus of activation may be necessary to ensure that the desired target is retrieved (Blaxton & Neely, 1983; Neill & Westberry, 1987). Although such narrowing might be achieved by enhancing the activation of target items, both retrieval speed and retrieval accuracy might benefit if an inhibitory mechanism suppressed interfering alternatives (Neill & Westberry, 1987). Some authors have assigned broad importance to such inhibitory processes, emphasizing their roles in excluding highly active but irrelevant representations from consciousness (Neill, 1991) and in reducing the opportunity for errors in cognition and action (Neill, 1991; Norman & Shallice, 1980; Tipper, Baylis, & Gordon, 1987). Other authors have limited the role of inhibitory processes to enhancing the retrieval of weakly represented memory items (Carr & Dagenbach, 1990; Dagenbach, Carr, & Barnhardt, 1990). In either case, inhibitory processes might play some role in the resolution of competition among simultaneously active memory targets.

Research on both episodic and semantic retrieval supports the existence of such inhibitory processes. For instance, if retrieval entails the suppression of interfering competitors, both the rate and the probability with which those competitors are retrieved on a later test should be impaired. We refer to this predicted effect as *retrieval-induced forgetting* (after Anderson, Bjork, & Bjork, 1994). Retrieval-induced forgetting has been shown to occur on semantic memory retrieval tasks. For example, Blaxton and Neely (1983) compared how reading versus generating (i.e., retrieving from semantic memory) prime items differentially affects the accessibility of semantically related target items. On priming trials in the *read-prime* condition, participants read aloud a presented category name (e.g., FRUIT) and then read aloud a presented exemplar of that category (e.g., ORANGE). On priming trials in the *generate-prime* condition, participants read aloud a presented category name; they were then shown a single-letter cue (e.g., O) and were to generate aloud a category exemplar beginning with that letter. Target trials in both conditions were identical to priming trials in the generate-prime condition; that is, participants were to generate aloud an exemplar of the preceding category beginning with the displayed letter (e.g., A). Two within-subjects independent variables used in that study are also of interest. The first is the number of prime trials; participants read or generated one or four primes. The second is whether the prime(s) and target were from the same or different categories; unrelated prime trials provided a baseline against which the facilitatory and inhibitory effects of related primes could be measured. Blaxton and Neely noted that traditional spreading activation theories predict that either reading or generating related prime items should activate target exemplars, speeding their production. In fact, read-prime participants did show facilitation in producing the target in both the one- and four-prime conditions; however, generate-prime participants showed a different pattern. Generating a single prime exemplar from the same category speeded target pro-

duction relative to generating a single unrelated prime, but generating four related primes slowed target production. Blaxton and Neely suggested that this difference supported a modified spreading activation theory (e.g., Keele & Neill, 1978) in which inhibitory processes resolve competition among simultaneously active competitors. Thus, when retrieval cues do not uniquely specify the target (e.g., *O*), inhibitory mechanisms may be invoked to weaken interfering nontargets, causing the speed with which those nontargets are retrieved on later occasions to be impaired. Similar conclusions have been drawn by several authors across a number of different paradigms involving retrieval from semantic memory (Brown, Whiteman, Cattoi, & Bradley, 1985; Carr & Dagenbach, 1990; Dagenbach & Carr, 1994b; Dagenbach et al., 1990; Neill & Westberry, 1987).

Retrieval also appears to be a source of forgetting from long-term episodic memory, as illustrated in a study conducted by Anderson et al. (1994). Anderson et al. used what they referred to as a *retrieval-practice paradigm* to explore both the positive and negative effects of retrieval on the long-term accessibility of information in episodic memory. The paradigm involved three phases: a study phase, a retrieval-practice phase, and a final test phase. In the study phase, participants studied several categories, each composed of several exemplars in a category-exemplar format (e.g., FRUIT ORANGE and FRUIT BANANA). Participants then engaged in directed "retrieval practice" on half of the items from half of the studied categories (three items from each of four categories) by completing category-plus-exemplar stem cue tests (e.g., FRUIT OR _____). After a 20-min retention interval, participants were shown each category name and asked to recall any exemplars of that category they remembered having seen at any point in the experiment. Of interest on this final test was not the recall of the practiced exemplars (e.g., FRUIT ORANGE)—which were expected to be facilitated—but rather the recall of unpracticed exemplars of practiced categories (e.g., FRUIT BANANA). Performance was measured relative to the recall for corresponding items in an unpracticed baseline category. On the basis of traditional spreading activation models, performing retrieval practice on FRUIT ORANGE should facilitate unpracticed exemplars like FRUIT BANANA. However, if retrieval involves active inhibitory processes that discriminate targets in memory, performing retrieval practice on FRUIT ORANGE should impair later recall of its unpracticed but interfering competitor. Consistent with the inhibitory model, Anderson et al. found that retrieval practice significantly impaired recall of unpracticed exemplars: Whereas 25% more items were recalled in the practiced than in the baseline condition, such facilitation came at the cost of an 11% decrement in the recall of unpracticed exemplars of practiced categories relative to the baseline condition.

The findings of Blaxton and Neely (1983) and those of Anderson et al. (1994) illustrate that retrieval does involve mechanisms that have deleterious effects on the subsequent processing of related information. Whether an item is currently being generated from semantic memory or is being recalled from an earlier episode, that item's performance may be hindered if it interfered during previous retrieval trials. Although the inhibitory process producing these instances of impaired recall has never been specified in detail, the literature contains at least two suggestions. First, discrimination of memory targets might be

achieved by automatic lateral inhibition processes akin to those at work in both perceptual and motor systems (Blaxton & Neely, 1983; Carr & Dagenbach, 1990; Dagenbach et al., 1990; Martindale, 1981; Roediger & Neely, 1982; for related proposals, see Estes, 1972; Konorski, 1967; McClelland & Rumelhart, 1981; Walley & Weiden, 1973). Second, inhibition might be achieved by some variety of active suppression, the mechanisms of which remain to be specified (Anderson & Bjork, 1994; Keele & Neill, 1978; Neill & Westberry, 1987; Neumann, Cherau, Hood, & Steinnagel, 1993). In this article, we use the term *inhibition* in a general way to refer to any such process that deactivates the representation of an item in memory. For present purposes, it is sufficient to note that both the functional case and the empirical case for inhibitory mechanisms in retrieval appear strong and that mechanistic theories of inhibitory phenomena could, in principle, be developed (a point to which we return later). Given these observations, one wonders why the status of inhibitory processes is an issue. That concern is the subject of our next section.

Noninhibitory Accounts of the Evidence

As plausible as the previous arguments for inhibition may seem, there are good reasons for viewing claims of inhibition with a critical eye. The mere observation that retrieval impairs memory performance of related items is not unequivocal evidence that an inhibitory mechanism is at work. Several classes of theories can explain these apparent inhibitory phenomena without proposing inhibitory processes (see Anderson & Bjork, 1994, for a review). We focus here on three of these classes of theory: *occlusion*, *resource diffusion*, and *associative decrement*.

To understand how these three noninhibitory mechanisms can explain memory inhibition effects, consider the simplified network representations depicted in Figure 1. In these diagrams, the category FRUIT and its exemplars ORANGE and BANANA are shown as nodes connected by associative links. Each node has a value representing its current level of activation in long-term memory, and each link has a strength, or weight, representing the strength of association between the category and the exemplar. For simplicity, we assume that retrieving an item involves the spread of activation from the category node (i.e., the retrieval cue) to the exemplars, with the amount of activation spread determined by the strength of the associative link along which it spreads. When an exemplar's activation exceeds a certain threshold level, that exemplar has been retrieved.

The manner in which the occlusion, resource diffusion, and associative decrement theories explain the seemingly inhibitory effects demonstrated by Anderson et al. (1994) can be seen by comparing the weights and activations in the baseline-category condition (Figure 1a) with those in the practiced-category condition (Figures 1b, 1c, and 1d). In the baseline-category condition, the weights and activations of the exemplars are equal, showing that these exemplars are equally accessible. In the practiced-category condition, however, the weight on the FRUIT-ORANGE link has been incremented to reflect the result of retrieval practice; consequently, the activation of ORANGE has also increased, leading to heightened accessibility. The crucial variation among the models, then, is the manner in which they explain the decreased recall of BANANA, the unpracticed exem-

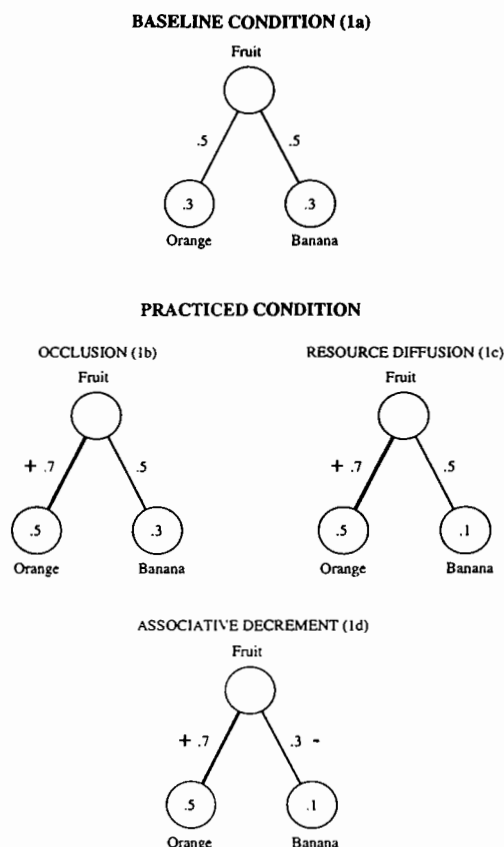


Figure 1. How noninhibitory models can account for the impairment of unpracticed items from practiced categories reported in Anderson, Bjork, and Bjork (1994). Circles represent nodes for categories and exemplars. Solid lines indicate associative links between categories and their exemplars; heavy lines and plus signs indicate links that were strengthened through retrieval practice of the exemplar. Numbers along the links indicate the strength of the link; numbers in the nodes indicate the activation of the exemplar.

plar. In the occlusion model (Figure 1b), the ability to recall BANANA decreases because the heightened accessibility of ORANGE leads ORANGE to block or occlude BANANA. For example, the greater weight of the link between FRUIT and ORANGE could allow it to reach retrieval threshold more quickly than BANANA, enabling the strengthened item to capture a limited-capacity response channel. The consequent intrusions of ORANGE when attempting to retrieve BANANA would ultimately lead participants to abandon the search for BANANA, lending the appearance that it has been inhibited when it has not. (Note the equivalence of the activation of BANANA in Figures 1a and 1b.) In the resource diffusion model (Figure 1c), the ability to recall BANANA decreases because the retrieval cue FRUIT no longer activates BANANA effectively. FRUIT spreads less activation to BANANA because this model assumes that activation is a finite resource to be distributed among the links emanating from a node as a function of the strength of the link; thus, strengthening the weight between FRUIT and ORANGE has the side effect of robbing activation that would otherwise have spread to BANANA. Although BANANA is less active as a consequence (compare Fig-

ure 1c and Figure 1a), this does not result from an inhibitory process acting on that item. Finally, in the associative decrement model (Figure 1d), the ability to recall BANANA decreases because the weight of the link between FRUIT and BANANA decreases during the practice of ORANGE. Weakening this association reduces FRUIT's capacity to activate BANANA, rendering that item less recallable. Again, although the critical item is less active as a consequence, this does not result from an inhibitory process acting on it.

These mechanisms clearly illustrate how memory effects that look like inhibition might not result from inhibitory processes at all; rather, such effects might merely reflect alternative non-inhibitory interference processes. These alternatives are not merely post hoc mechanisms designed to explain memory inhibition in a noninhibitory way; in fact, they each form components of larger architectures capable of explaining a broad array of phenomena. For example, occlusion processes have enjoyed a long-standing popularity as a means of explaining memory interference effects (McGeoch, 1936; Melton & Irwin, 1940; Mensink & Raaijmakers, 1988; Rundus, 1973), perhaps motivating the relative-strength assumptions that form the bedrock of some theories of memory retrieval (e.g., Gillund & Shiffrin, 1984; Mensink & Raaijmakers, 1988; Raaijmakers & Shiffrin, 1981). Resource diffusion finds its conceptual heritage in limited resource theories of attention (e.g., see Broadbent, 1958; Kahneman, 1973; Norman & Bobrow, 1975; Wickens, 1984); its extension to memory phenomena has been most clearly articulated in J. R. Anderson's (1983) Adaptive Control of Thought (ACT*) cognitive architecture, forming the basis of that model's general approach to spreading activation. Finally, associative decrement hypotheses in memory research date back to Melton and Irwin's (1940) two-factor theory of interference, in which such processes were thought to account for the empirical phenomena of unlearning. Although associative decrement approaches to memory interference fell from favor (see, e.g., Postman, Stark, & Fraser, 1968), they have acquired renewed force with the application of modern connectionist learning algorithms to these phenomena (Lewandowsky, 1991; Sloman & Rumelhart, 1992). Thus, these alternative accounts of memory inhibition phenomena originate from larger theoretical frameworks, the ability of which to explain a broad array of findings renders them especially difficult to ignore.

Empirical Ambiguities

It is clear from the preceding discussion that the case for inhibitory processes in memory retrieval cannot rest solely on the observation that retrieval processes impair the recall for competing memory items; rather, the case for inhibitory processes in retrieval requires empirical studies directed specifically at resolving the ambiguity created by noninhibitory interference theories. The few attempts that have been made to distinguish these approaches have favored an inhibitory interpretation. For instance, several investigations have shown that strengthening category exemplars impairs competing items only when the process of strengthening requires the discrimination of those items from interfering memory competitors (e.g., as in retrieval but not mere presentation; see Anderson, Bjork, & Bjork, 1993; Anderson et al., 1994; Blaxton & Neely, 1983); contrary to the

predictions of the noninhibitory theories discussed in the previous section, the strengthening of a memory item, by itself, often does not impair competing exemplars.¹

Although these findings support inhibitory theories of retrieval, it is difficult to entirely exclude two classes of alternative explanations for the failure of strengthening to produce impairment. First, the amount by which items are strengthened may be questioned; if competing items are not sufficiently strengthened, then occlusion or resource diffusion effects might not be observable (see, e.g., Blaxton & Neely, 1983). Second, additional facilitatory processes acting on competing items (e.g., participant rehearsal strategies or concurrent spreading activation) might be proposed to offset the impairment expected on the basis of noninhibitory theories. The flexibility of such counterarguments can render even the most convincing evidence that inhibitory mechanisms underlie retrieval-induced forgetting consistent with noninhibitory theories. Some authors have concluded that, without further constraints on both inhibitory and noninhibitory theories, there may be no principled way to distinguish among them (Blaxton & Neely, 1983; Roediger & Neely, 1982). Thus, the status of inhibitory processes in memory retrieval remains a point of debate.

RESOLVING THE DEBATE: DISTINGUISHING INHIBITORY AND NONINHIBITORY THEORIES OF RETRIEVAL

Another way to distinguish inhibitory and noninhibitory theories is to identify a situation in which only inhibitory approaches predict impairment. Such a circumstance emerges when one considers a property shared by all of the noninhibitory theories discussed thus far: All noninhibitory theories predict that retrieval-induced forgetting should be restricted to only those competitors that were also studied and tested under the retrieval-practice cue. Consider the role of the retrieval cue FRUIT in the previous illustrations: (a) In occlusion models, the recall of BANANA decreases because the cue FRUIT persistently elicits the strengthened competitor, ORANGE, blocking BANANA; (b) in resource diffusion models, the recall of BANANA decreases because activation spreading from the cue FRUIT is diverted to the strengthened item ORANGE; and (c) in associative decrement models, the recall of BANANA decreases because its link to the cue FRUIT is weakened during strengthening of the competing association between FRUIT and ORANGE. In each of these examples, impairment derives from the relationship of the critical item BANANA to the retrieval cue FRUIT; items that are studied and tested under cues other than FRUIT ought not to suffer from strengthening the FRUIT-ORANGE association. In contrast, inhibitory theories predict impairment for items studied and tested under distinct retrieval cues: As long as one item interferes during the retrieval practice of another, whether the two items share a cue at test should be irrelevant.

This observation suggests an adaptation of the retrieval-practice paradigm of Anderson et al. (1994) that allows the case for inhibitory processes to be critically examined. Consider the example provided in Figure 2.

In this example, participants have studied two categories: RED items and FOOD items. The RED category contains the exemplars BLOOD and TOMATO; the FOOD category contains the




Figure 2. Construction of related categories. Solid lines indicate studied category-exemplar pairs; heavy lines indicate the subset of those pairs that received retrieval practice; dashed lines indicate a preexisting category-exemplar relation that was not studied in the experiment.

examples STRAWBERRY and CRACKERS. Note that the two center items, TOMATO and STRAWBERRY, although studied under only one of the categories, are also semantic members of the other category. This implicit semantic connection is represented by the dashed lines. Suppose that participants have performed retrieval practice on the item RED-BLOOD, as denoted by the darkened line connecting the category and exemplar. The previous findings of Anderson et al. (1994) and both the inhibitory and noninhibitory theories of retrieval discussed thus far predict that retrieval practice of RED-BLOOD should impair the later recall of RED-TOMATO.

The crucial question in this example concerns the fate of the item FOOD-STRAWBERRY. Assuming that STRAWBERRY is both studied and tested under the retrieval cue FOOD, all of the noninhibitory theories predict that the recall of STRAWBERRY should not be impaired by the retrieval practice of RED-BLOOD. BLOOD will not occlude STRAWBERRY because the strengthened item BLOOD is not associated with the retrieval cue FOOD; for similar reasons, none of STRAWBERRY's activation resources are robbed by the strengthening of RED-BLOOD. Finally, the FOOD-STRAWBERRY association should not be decremented by the strengthening of RED-BLOOD because it is not a competing association. The inhibitory account of the Anderson et al. (1994) impairment, however, suggests that practicing RED-BLOOD might impair later recall of FOOD-STRAWBERRY. This might occur because the implicit semantic link between RED and STRAWBERRY should activate STRAWBERRY during the retrieval practice of RED-BLOOD, causing it to interfere with the ability to discriminate BLOOD. If STRAWBERRY does interfere with the practice of BLOOD, then STRAWBERRY should be just as susceptible to inhibition as those items nested explicitly under

¹ Although several well-known findings suggest that strengthening items impairs competitors, the case that the strengthening causes the impairment is weak. For instance, increases in retroactive interference with the degree of interpolated learning (Barnes & Underwood, 1959), increases in part-set cuing inhibition with increasing numbers of part-set cues (Rundus, 1973), and the list-strength effect (Ratcliff, Clark, & Shiffrin, 1990) may all be cited as evidence that strengthening impairs recall of competitors. Anderson, Bjork, & Bjork (1994) have argued, however, that each of these findings either confounds strengthening with some form of retrieval-induced forgetting (e.g., output interference) or does not, in fact, show the expected impairment. When retrieval-induced forgetting is controlled, strengthening frequently fails to impair competitors.

the retrieval cue RED. Thus, testing with an independent retrieval cue should enable one to empirically distinguish inhibitory and noninhibitory theories of retrieval-induced forgetting. This method, which we call the independent probe technique, is the basis for all of our experiments directed at assessing the status of cognitive inhibition.

GENERAL METHOD: INDEPENDENT PROBE TECHNIQUE

The procedure and design for most of the experiments reported here were similar. Thus, we describe the general method now, specifying special features and exceptions as we introduce each experiment.

The method used here adapted the retrieval-practice procedure of Anderson et al. (1994) to the *independent probe technique*. As in the Anderson et al. studies, there were three phases: a learning phase, a retrieval-practice phase, and a final test phase. In the learning phase, participants studied several six-exemplar categories in category-exemplar paired-associate formats (e.g., RED-BLOOD). In the retrieval-practice phase, half of the items from half of the categories received retrieval practice by means of cues such as RED BL_____. In this phase, participants were instructed to retrieve the exemplar that began with the specified letters and was paired with the category in the earlier learning phase. After a 20-min retention interval, participants were shown each category name in turn and given 30 s per category to recall as many studied exemplars as they could. This procedure created two types of categories: those that involved retrieval practice (hereafter Rp categories) and those that did not (hereafter Nrp categories). Within the practiced categories, there were two types of exemplars: those that were actually practiced in the retrieval-practice phase (hereafter Rp+ items) and those that were not (hereafter Rp- items). To demonstrate the facilitatory effects of retrieval practice, we compared final test performance for Rp+ (practiced) items with performance on those same items when no items from the category were practiced (i.e., an Nrp baseline); to demonstrate retrieval-induced forgetting, we also compared performance on Rp- items with performance on those same items when they appeared in the Nrp baseline condition.

In the Anderson et al. (1994) retrieval-practice procedure, categories were constructed to be as dissimilar or unrelated to one another as possible. As suggested by the example given in Figure 2, however, adapting the retrieval-practice procedure to the independent probe technique required the construction of pairs of related categories that contained similar items. Each new category contained six exemplars. Three of these six items, in addition to being members of their own study category, were also similar to members of the related category (*similar* items), and the remaining three items were not (*dissimilar* items). (Appendixes A and B show two different ways in which the items were similar across our experiments.) For example, the related categories depicted in Figure 3—RED and FOOD—each contained six items, some of which were similar by virtue of their being semantically related to both studied categories (e.g., TOMATO and STRAWBERRY) and some of which were not similar by that criterion (e.g., BLOOD and CRACKERS). Construction of these pairs of categories allowed us to have a condition in which

Independent Probe Technique

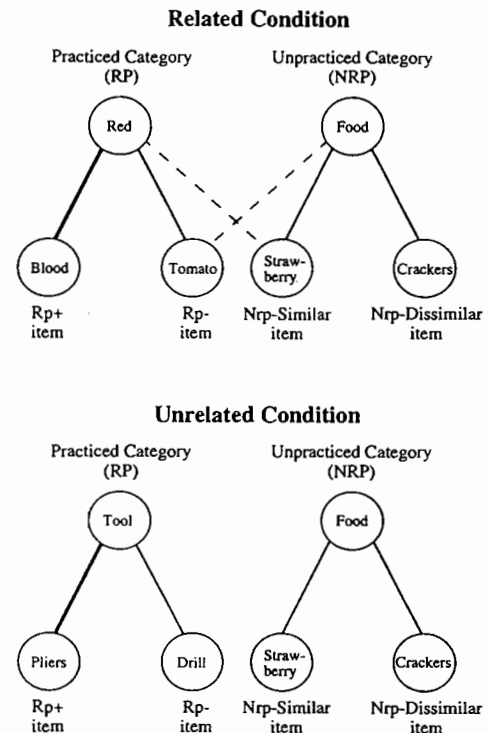


Figure 3. General design of the present experiments and the critical comparison using the independent probe technique. Solid lines indicate studied category-exemplar pairs; heavy lines indicate the subset of those pairs that received retrieval practice; thin dashed lines indicate a preexisting category-exemplar relation that was not studied in the experiment. The shaded circles show the critical comparison in most of the experiments reported in this article.

participants studied a pair of related categories (the *related* condition) and a condition (replicating Anderson et al.) in which they studied two categories that were not related (the *unrelated* condition; see Figure 3). The new related condition enabled us to ask the novel question of this design: What effect does performing retrieval practice on items like RED-BLOOD have on the later recall of an item such as FOOD-STRAWBERRY, which is semantically similar to it but was studied under another cue? As argued earlier, inhibitory theories of retrieval-induced forgetting predict that the recall of STRAWBERRY should be impaired by the practice of RED-BLOOD; that is, they predict *cross-category inhibition*.

The standard within-category retrieval-induced forgetting effect was measured in the unrelated condition by comparing Rp- and Nrp items; cross-category impairment, however, had to be measured in the novel way depicted in Figure 3. Note that, in this example, the critical item STRAWBERRY is what we call an *Nrp-Similar* item, because it is an exemplar from an Nrp category (FOOD) and because it is similar to items from another category in the experiment. In contrast, CRACKERS, also a member of the FOOD category, is an *Nrp-Dissimilar* item because it is not similar to items from another category in the

experiment by the criterion just mentioned. To assess impairment of the critical Nrp-Similar items, we compared their recall performance in the related condition (in which the words they were similar to were also studied by the participant) with their recall performance in the unrelated condition (in which the words they were similar to were not studied by the participant; see the shaded circles in Figure 3). Recall of Nrp-Similar items in the unrelated condition provided an appropriate baseline because we did not expect those items to interfere with the practice of items from an unrelated category. For instance, we did not expect the critical item FOOD-STRAWBERRY to be impaired by the retrieval practice of TOOL-PLIERS because STRAWBERRY was unrelated to the TOOL category and thus should not have interfered with the practice of TOOL-PLIERS. Thus, if Nrp-Similar items were recalled less well in the related condition than in the unrelated condition, we would have evidence for cross-category inhibition.

EXPERIMENT 1

To determine whether retrieving an item from long-term episodic memory impairs subsequent recall of items studied and tested under a different category cue, we used the independent probe methodology described previously. We constructed three pairs of six-exemplar categories; each category contained three exemplars that could be classified under the related category and three that could not (see Appendix A). Of these six categories, a given participant studied only four: a pair of categories in the related condition and two unrelated categories in the unrelated condition (see Table 1). Participants performed retrieval practice on the dissimilar items from one related category and one unrelated category, and we tested their recall for all items 20 min later with a category-cued free recall test.

Regardless of whether inhibitory mechanisms underlie retrieval-induced forgetting, the present experiment should replicate the basic within-category retrieval-induced forgetting observed by Anderson et al. (1994); that is, in the unrelated condition, participants' recall performance should be worse for Rp- items than for Nrp-Similar items. The crucial question, however, centers around the fate of Nrp-Similar items in the related condition versus the unrelated condition. If noninhibitory mechanisms such as occlusion, resource diffusion, and associative decrement are the sole source of the impairment of Rp- items, then Nrp-Similar items in the related condition should remain unaffected by retrieval practice of Rp+ items. Because participants studied and retrieved those items under an independent category retrieval cue, strengthening Rp+ items should not block Nrp-Similar items, rob their activation resources, or cause decreases in their strengths of association to the Nrp category. Thus, Nrp-Similar recall performance should not differ across the related and unrelated conditions.

If, however, inhibitory processes cause retrieval-induced forgetting, retrieval practice should impair both Rp- and Nrp-Similar items in the related condition; because their prior semantic relation to the practiced category might cause these items to interfere during retrieval practice of Rp+ items, they should be inhibited regardless of whether participants studied and retrieved those items under the same cue as or a different cue than practiced items. If such interference does occur then,

as described earlier, final recall performance for Nrp-Similar items in the related condition should be impaired relative to final recall performance for Nrp-Similar items in the unrelated condition, providing evidence for cross-category inhibition.

Method

Participants

Forty-eight students from an introductory psychology class at the University of California, Los Angeles, participated in partial fulfillment of a course requirement.

Design

The experiment involved a 2×4 within-subjects design. First, we manipulated the relatedness of category pairs. Two categories appeared in the related condition; each of those categories contained three items that were semantically members only of that category and three items that were semantically members of both related categories. Two categories appeared in the unrelated condition; each contained six items that were semantically members only of that studied category. The categories and words are shown in Appendix A. Relatedness of category pairs was manipulated in the study phase of the experiment by varying the combinations of categories studied by participants; Table 1 illustrates the combinations of categories studied by various participants in the experiment.

Second, we manipulated the retrieval-practice status of an exemplar on four levels: Rp+, Rp-, Nrp-Dissimilar, and Nrp-Similar. As in the Anderson et al. (1994) procedure, participants performed retrieval practice on half of the studied categories; thus, for a given participant, two of the four studied categories were practiced (Rp) categories; and two were unpracticed (Nrp) categories. For each Rp category, participants practiced three exemplars (Rp+ items) and did not practice the remaining three (Rp- items). For example, as illustrated in Table 1 for Participant 1, some participants practiced the exemplars RED-BLOOD, RED-FIRE, and RED-SUNBURN from the Rp category RED, making those exemplars Rp+ items; the remaining three unpracticed exemplars from that category, RED-TOMATO, RED-APPLE, and RED-CHERRY, then became Rp- items. For each Nrp category, three of the six exemplars were similar to other items in the experiment (Nrp-Similar items) and three were dissimilar to other items in the experiment (Nrp-Dissimilar items). For example, again as shown for Participant 1 in Table 1, some participants studied the following exemplars of the Nrp category FOOD: STRAWBERRY, RADISH, KETCHUP, CRACKERS, PEAS, and BREAD. The first three exemplars would be Nrp-Similar items because they are also semantically red things, but the second three exemplars would be Nrp-Dissimilar items because they are not also red things. Note that, for Participant 1, the Nrp category FOOD was in the related condition, so the Nrp-Similar items were similar to other items actually studied by the participant (i.e., some of the foods were also red). For an Nrp category in the unrelated condition (e.g., LOUD for Participant 1), Nrp-Similar items (LAWNMOWER, SANDBLASTER, and COMPRESSOR) were similar to other items we constructed (i.e., the tools FILE, PLIERS, and SCREWDRIVER), but those other items were not studied by the participant.

As illustrated by the preceding examples, and as shown in Table 1, the Rp+ exemplars in this experiment were always dissimilar items, and the Rp- items were similar items. Note that because participants practiced only dissimilar items, the sets of Rp+ and Nrp-Dissimilar items were identical, as were the sets of Rp- and Nrp-Similar items, after consideration of all of our counterbalancing measures. We recorded the percentage of items correctly recalled in the final category-cued recall test.

Table 1
Illustration of Learning-Practice Combinations for Half of the Participants in Experiment 1

Category	Learning Booklet 1		Learning Booklet 2		Learning Booklet 3	
	Participant 1	Participant 2	Participant 3	Participant 4	Participant 5	Participant 6
Red	Blood Fire Sunburn Apple Cherry Tomato	Blood Fire Sunburn Apple Cherry Tomato	Blood Fire Sunburn Apple Cherry Tomato	Blood Fire Sunburn Apple Cherry Tomato		
Food	Bread Crackers Peas Ketchup Radish Strawberry	Bread Crackers Peas Ketchup Radish Strawberry			Bread Crackers Peas Ketchup Radish Strawberry	Bread Crackers Peas Ketchup Radish Strawberry
Fly	Kite Glider Frisbee Butterfly Eagle Ladybug	Kite Glider Frisbee Butterfly Eagle Ladybug	Kite Glider Frisbee Butterfly Eagle Ladybug	Kite Glider Frisbee Butterfly Eagle Ladybug		
Animal			Sheep Giraffe Hamster Wasp Bat Pigeon	Sheep Giraffe Hamster Wasp Bat Pigeon	Sheep Giraffe Hamster Wasp Bat Pigeon	Sheep Giraffe Hamster Wasp Bat Pigeon
Loud	Thunder Yell Traffic Lawnmower Sandblaster Compressor	Thunder Yell Traffic Lawnmower Sandblaster Compressor			Thunder Yell Traffic Lawnmower Sandblaster Compressor	Thunder Yell Traffic Lawnmower Sandblaster Compressor
Tool			File Pliers Screwdriver Jackhammer Drill Chainsaw	File Pliers Screwdriver Jackhammer Drill Chainsaw	File Pliers Screwdriver Jackhammer Drill Chainsaw	File Pliers Screwdriver Jackhammer Drill Chainsaw

Note. Related pairs of categories are vertically adjacent. Each participant studied items down the column. Practiced items appear in boldface. Within each category, the first three items are dissimilar items; the last three are similar items. Six additional participants would be necessary to show the complete design; for them, the similar items would be exchanged between related categories.

Procedure

We collected data from the participants individually or in groups of up to 10. They were told that they would be participating in an experiment on memory and reasoning and that the experiment would have several parts. The experiment was conducted in four phases: a learning phase, a retrieval-practice phase, a distractor phase, and a surprise category-cued recall phase.

Learning Phase

Participants were given a learning booklet and were told that (a) each page in the booklet contained a category-example pair and they were to try to remember the category examples as best they could; (b) they would be given 5 s to study each category-example pair and they were supposed to spend all of this time relating the example to its category; (c) at the end of 5 s, a voice on a tape recording would instruct them to turn the page; and (d) this sequence was to be repeated until all pages in the learning booklet had been studied.

Retrieval-Practice Phase

When participants completed the learning phase, the experimenter collected learning booklets and distributed practice booklets. Participants were told that (a) each page in the new booklet contained one of the category names that they had seen in the previous phase of the experiment and the first two letters of one of the members of that category, (b) they would be given 10 s to try to remember the example and to write it down on the page, (c) they were to retrieve something that they had already seen rather than simply figure out any word that fit, and (d) some of the examples might be tested more than once, in which case they were supposed to do the same thing each time. As in the learning phase, a tape-recorded voice instructed them when to turn pages.

Distractor Phase

After the retrieval practice phase had been completed, booklets were collected, and participants were given an unrelated reasoning task that took 20 min.

Final Test Phase

Participants were given recall booklets and were told that (a) at the top of each page, there would be the name of one of the categories studied previously; (b) they would be given 30 s to try to recall all examples of that category that they had seen at any time in the experiment; and (c) after 30 s, the experimenter would instruct them to turn to the next page. When participants completed this task, they were debriefed and dismissed.

Materials

Categories

We constructed 10 categories, 4 filler and 6 experimental. The 6 experimental categories consisted of three pairs of related categories (Appendix A provides the complete list of items for all three of these pairs of categories). Each category of a pair contained three items that were uniquely members of that particular category (i.e., dissimilar items). Across the related pair of categories, however, there were six items that semantically were members of both categories (i.e., similar items). For example, the categories RED and FOOD formed one pair of related categories. RED included three dissimilar items such as BLOOD; FOOD included three dissimilar items such as CRACKERS. In addition, there were six items such as TOMATO and STRAWBERRY that were similar items (three for each category). Note that similar items could be categorized under either category of the related pair; dissimilar items could not. The filler categories—mountains, cities, countries, and composers—consisted of six items each.

Category members were selected according to the following constraints. First, we required that exemplars be unambiguous members of their categories. For example, we rejected the item DRESS as an exemplar of the RED category because, although dresses are often red, they are not necessarily red, nor are dresses red by default. We selected the item BLOOD as a member of the RED category, however, because most people assume blood to be red and red is strongly associated with the word BLOOD. Second, we excluded exemplars if they had strong associations with other exemplars that had already been selected, to minimize participants' ability to use interexemplar retrieval strategies during the final test phase. Third, to ensure that each exemplar's two-letter cue in the retrieval-practice task uniquely designated that exemplar, we did not allow two words in the experiment to begin with the same first two letters. Finally, to avoid interference from extra-experimental items, we did not allow any chosen category member to have the same first two letters as a familiar unchosen category member. For example, we excluded the item DOG as an ANIMAL because participants might intrude DONKEY.

Learning Booklets

Learning booklets consisted of 48 pages, each containing the name of a category and an exemplar of that category to its right (e.g., RED BLOOD), centered on the page. Each booklet contained all six exemplars from the four filler categories, as well as six exemplars—three dissimilar and three similar—from each of four of our six experimental categories. The four experimental categories always consisted of one of the pairs of related categories and one category from each of the remaining two related pairs, implementing our related and unrelated conditions, respectively (see Table 1). We ordered these categories and exemplars in the learning booklet in such a way as to avoid the formation of both within-category and between-categories associations. Categories were prohibited from appearing twice in a row, and related categories were prohibited from occurring adjacently. In addition, we generally prevented unrelated categories from appearing next to each other more than once. To control for serial position effects, we ensured that the

average study position for each category (and, therefore, for the unrelated vs. related and the to-be-practiced vs. not-to-be-practiced categories) was approximately equal. These restrictions applied to the experimental categories; filler categories were used liberally to enforce the constraints for the experimental categories. The first 3 and last 3 pages of the learning booklets contained filler category items to control for primacy and recency effects.

Each of our categories appeared equally often in the related and unrelated conditions. In addition, our critical similar exemplars appeared equally often in both related categories (e.g., TOMATO was studied half of the time under FOOD and half under RED). The former but not latter manipulation is illustrated in Table 1.

Retrieval-Practice Booklets

Retrieval-practice booklets consisted of 51 pages, each containing the name of a studied category and the first two letters of a studied member of that category followed by a solid blank line (e.g., RED BL_____), centered on the page. We held the length of the blank line constant so as to give no cues for word length.

We constructed retrieval-practice booklets as follows. Each booklet tested participants on six categories: four filler and two critical experimental categories. The two experimental categories always included one related and one unrelated category. From these two categories, we selected only the dissimilar items for the retrieval-practice booklet. Table 1 shows practiced items for each participant in boldface. We constructed the practice booklets by inserting each critical exemplar into the booklet three times in an expanding sequence. The first and second practices occurred with an average of 3.0 intervening items; the second and third practices occurred with an average of 6.7 intervening items. To reduce opportunities for interassociation among exemplars, we prohibited members of the same category from appearing adjacently in the booklet, and, in general, we did not allow items to be practiced adjacently more than once. To minimize order effects, we equated the average practice position for items in the two practiced categories. These restrictions applied to the experimental categories; again, filler category items were used liberally to enforce the previous constraints. Filler items always occupied the first and last three practice positions in the booklet.

Each of our experimental categories was practiced and unpracticed equally often (compare Participants 1 and 2 in Table 1, who learned the same four categories but practiced two complementary categories). In addition, to ensure that the particular way in which we inserted exemplars into the practice booklets had no impact on our results, we created two different random practice orders for each set of practiced items.

Recall Booklets

Recall booklets consisted of four pages, each containing the name of a learned experimental category at the top. We constructed the final recall booklets as follows. The first page always contained the name of a category in the related condition, the second page contained the name of a category in the unrelated condition, the third page contained the name of the other related category, and the fourth page contained the name of the other unrelated category. For half of the participants, the practiced categories were recalled first and last; for the other half of the participants, the practiced categories were recalled second and third (as a result of retrieval-practice counterbalancing). Thus, related categories were never recalled adjacently. Across all participants, the average output position of the practiced and unpracticed categories was the same. Different test booklets were constructed to correspond with each learning booklet so that participants recalled items only from categories they had studied. There were two different recall orders for each learning booklet.

Results and Discussion

All analyses were performed with learning booklet and practice counterbalancing as between-subjects conditions.

Retrieval Practice

We counted whether participants successfully retrieved the correct item during each retrieval practice trial for the experimental (not filler) categories. Because participants practiced six experimental items (three from one of the categories in the related condition and three from one of the categories in the unrelated condition) each three times, there were 18 relevant practice trials. The retrieval-practice success rate for items from the category in the unrelated condition (69%) did not differ significantly from the retrieval-practice success rate for items from the category in the related condition (73%), $F(1, 42) = 1.04$, *ns*.

Final Recall

Note, for purposes of comparing recall percentages, that in this design Rp+ and Nrp-Dissimilar items were identical, as were Rp- and Nrp-Similar items (see Table 2). We report the

Table 2
Experiments 1, 2, and 3: Mean Percentage of Words Recalled as a Function of Category Relatedness and Retrieval-Practice Status of Items

Category relatedness	Retrieval-practice status			
	Practiced category (RED)		Unpracticed category (FOOD)	
	Rp+ (BLOOD)	Rp- (TOMATO)	NrpS (STRAWBERRY)	NrpD (CRACKERS)
Experiment 1 (N = 48)				
Unrelated	69	24	38	44
Related	74	22	22	36
Category relatedness	Practiced category (GREEN)		Unpracticed category (SOUPS)	
	Rp+ (EMERALD)	Rp- (LETTUCE)	NrpS (MUSHROOM)	NrpD (CHICKEN)
	Experiment 2 (N = 54)			
Unrelated	72	42	52	43
Related	65	36	37	43
Experiment 3a (N = 27)				
Unrelated			48	36
Related			48	44

Note. Examples of items in each condition are shown in parentheses; they correspond to the examples in Figures 2 and 4. Unrelated categories were learned without their paired category in the learning booklet; related categories were learned with their paired category in the learning booklet. The comparisons for detecting cross-category inhibition for these experiments are shown in boxes. Rp+ = practiced items; Rp- = unpracticed items; NrpS = similar items; NrpD = dissimilar items. Rp+ and NrpD are the same item set; Rp- and NrpS are the same item set.

results of planned comparisons based on percentage correct measures. First, to verify that the retrieval-practice procedure had the expected positive and negative effects on Rp+ and Rp- items, respectively, we restricted comparisons to the unrelated condition. The materials in the unrelated condition resembled those of Anderson et al. (1994), in which the Rp and Nrp categories had no between-categories similarity. As expected, retrieval practice significantly facilitated the later recall of Rp+ items ($M = 69\%$) relative to Nrp-Dissimilar baseline items ($M = 44\%$), $F(1, 36) = 35.68$, $p < .001$, whereas it impaired recall of Rp- items ($M = 24\%$) relative to Nrp-Similar baseline items ($M = 38\%$), $F(1, 36) = 8.40$, $p < .01$. These findings replicate both the retrieval-based learning and retrieval-induced forgetting effects observed by Anderson et al.

Support for Inhibitory Mechanisms

The novel question in the present design concerned the effect of retrieval practice on Nrp-Similar items in the related condition. Planned comparisons revealed that retrieval practice impaired the recall of Nrp-Similar items in the related condition ($M = 22\%$) relative to performance of those same items in the unrelated condition ($M = 38\%$; see the boxed items in the upper part of Table 2), $F(1, 36) = 8.36$, $p < .01$. That is, performing retrieval practice on Rp+ items (e.g., RED-BLOOD) impaired both Rp- items (e.g., RED-TOMATO) and Nrp-Similar items (e.g., FOOD-STRAWBERRY), even though, in the latter case, participants encoded and retrieved items under a different category cue. Thus, cross-category inhibition occurred.

The present findings therefore favor the inhibitory account of retrieval-induced forgetting over all of the noninhibitory theories discussed previously (and over several others as well; see Anderson & Bjork, 1994). This conclusion follows because none of the processes invoked by those models to account for within-category impairment pertain to the recall of items using a separate retrieval cue. Although the strengthening of the associative link between RED and BLOOD may occlude, divert resources from, or decrease the association with the item TOMATO, that strengthening should be entirely irrelevant to the recall of STRAWBERRY given the encoding and retrieval cue FOOD. According to the inhibition model of retrieval, however, repeated retrieval of RED-BLOOD should require the inhibition of all items that interfere with the retrieval process; just as items studied under the same retrieval cue are likely to interfere, items such as FOOD-STRAWBERRY, which have prior semantic links to RED, may interfere as well. Because inhibition affects the item itself and not an association, impairment of such items should be observable even when measuring recall with the independent retrieval cue FOOD. Thus, noninhibitory mechanisms cannot explain the impaired recall of Nrp-Similar items, so it appears necessary to posit inhibitory mechanisms in directed memory retrieval.

Does Encoding Specificity Protect Against Inhibition?

The magnitude of the impairment for Nrp-Similar items (Nrp-Similar-unrelated items - Nrp-Similar-related items = 16%) is especially surprising when one considers the encoding specificity principle (Tulving & Thomson, 1973). This princi-

ple asserts that the likelihood of recalling an item hinges on the match between cues at encoding and those available at retrieval. For example, if, after studying a list containing the paired associate GLUE-CHAIR, participants' recall is tested with the novel cue TABLE_____ (a high associate of CHAIR), performance is far worse than if memory is tested with the original cue GLUE (a low associate of CHAIR). In examples such as this, the decreased ability of TABLE to elicit CHAIR is often explained by assuming that the encoding cue GLUE biases the representation of CHAIR so that TABLE is no longer a potent cue. In the context of the present experiment, then, STRAWBERRY's encoding as a member of the FOOD category should have diminished participants' tendency to spontaneously generate that item during retrieval practice of RED-BLOOD. Consequently, STRAWBERRY should have remained unimpaired or suffered less impairment than Rp- items encoded explicitly under the practiced category, irrespective of whether inhibitory mechanisms existed. In fact, the impairment of Nrp-Similar items (16%) appeared comparable to the impairment of Rp- items (Nrp-Similar-unrelated items - Rp- unrelated items = 14%). This finding suggests the intriguing possibility that encoding events in distinct contexts may not necessarily protect such events from retrieval-based inhibition.

Reporting Bias as a Possible Explanation

It is possible that the cross-category impairment was an artifact accidentally produced by our related condition. Perhaps participants in the related condition recalled Nrp-Similar items perfectly well but then became confused as to which category those items belonged. For example, participants cued with the category FOOD on the final test might have recalled STRAWBERRY but then become confused about whether they had studied STRAWBERRY with FOOD or RED. Such confusion may have led them to withhold responses for Nrp-Similar items in the related condition. Because similar confusion about category membership could not have occurred in the unrelated condition, reporting bias might have caused the apparent inhibitory effect for Nrp-Similar items. Intrusion data collected in the present experiment lend support to this hypothesis. In the final recall phase, participants intruded 48 items out of a total of 518 responses (an intrusion rate of 9.6%). More intrusions occurred in the related than in the unrelated condition (33 items vs. 15 items), and half of the intrusions in the related condition (16 items) resulted from participants recalling a studied item to the wrong category cue. If confusion over the category membership of Nrp-Similar exemplars lowered recall scores, our results would not constitute evidence for inhibitory processes.

Unanticipated Results

One unexpected feature of the present findings was the impaired recall performance for Nrp-Dissimilar items in the related condition ($M = 36\%$) relative to the performance on those same items in the unrelated condition ($M = 44\%$), $F(1, 36) = 4.00$, $p = .05$. We did not expect this impairment because we designed Nrp-Dissimilar items specifically so that they bore no resemblance to Rp+ items and were not semantic members of the Rp category. For example, the Nrp-Dissimilar item FOOD-

CRACKERS does not resemble the Rp+ item RED-BLOOD, nor are actual crackers usually red. Therefore, there is no straightforward basis for expecting CRACKERS to interfere during retrieval practice of RED-BLOOD; thus, CRACKERS should not be vulnerable to inhibition. Nonetheless, practice on RED-BLOOD impaired FOOD-CRACKERS.

An accidental but systematic similarity between Nrp-Dissimilar and Rp- items may have contributed to this unexpected impairment. Although Nrp-Dissimilar items were indeed dissimilar to Rp+ items and were not members of the Rp category, they did share a category with Rp- items. Figure 2 illustrates this point: The Nrp-Dissimilar item FOOD-CRACKERS resembles RED-TOMATO, the Rp- item, in that semantically they are both foods. It is possible that merely bearing similarity to an item that is directly vulnerable to inhibition, such as RED-TOMATO, renders the similar item vulnerable by mechanisms we did not anticipate. Experiment 2 addressed this issue, as well as the issue of reporting bias.

EXPERIMENT 2

Experiment 1 demonstrated that performing retrieval practice on certain category exemplars can impair the later recall of similar items studied and tested under different category cues. These findings contradict the predictions of all three of the non-inhibitory theories reviewed earlier, which assert that impairment stems from competition among items sharing a common retrieval cue. These findings also appear to provide striking support for the notion that inhibitory processes assist memory retrieval.

As we noted, however, the present evidence for cross-category inhibition might have arisen as an artifact of our category relatedness manipulation: Participants might have recalled Nrp-Similar items but then failed to report them because they were uncertain as to which category the items belonged. To evaluate this reporting-bias hypothesis, Experiment 2 used new categories that eliminated uncertainty regarding the category membership of Nrp-Similar items. Figure 4 depicts the relationship between these new categories.

In this example, participants study the categories GREEN and SOUPS and then perform retrieval practice on the GREEN category. In particular, participants practice the item GREEN-EMERALD (an Rp+ item), making GREEN-LETTUCE its Rp- competitor, SOUPS-MUSHROOM the Nrp-Similar item, and SOUPS-CLAM the Nrp-Dissimilar item. This arrangement of categories and items resembles that of Experiment 1 in that there are certain critical items in the Nrp category that are similar to items in the Rp category. The novel feature of these categories and exemplars, however, centers around how Nrp-Similar items are related to the Rp category. In Experiment 1, Nrp-Similar items were similar to the Rp category because those items were semantically members of the Rp category (e.g., STRAWBERRY, studied as a FOOD, also happens to be red, and RED was an Rp category). In the materials for Experiment 2, Nrp-Similar items were not members of the Rp category; rather, their similarity to the Rp category rested solely on their similarity to Rp- items. For example, the Nrp-Similar item SOUPS-MUSHROOM could not be mistaken as a member of the Rp category GREEN because mushrooms, in general, are not green; however, MUSHROOM is

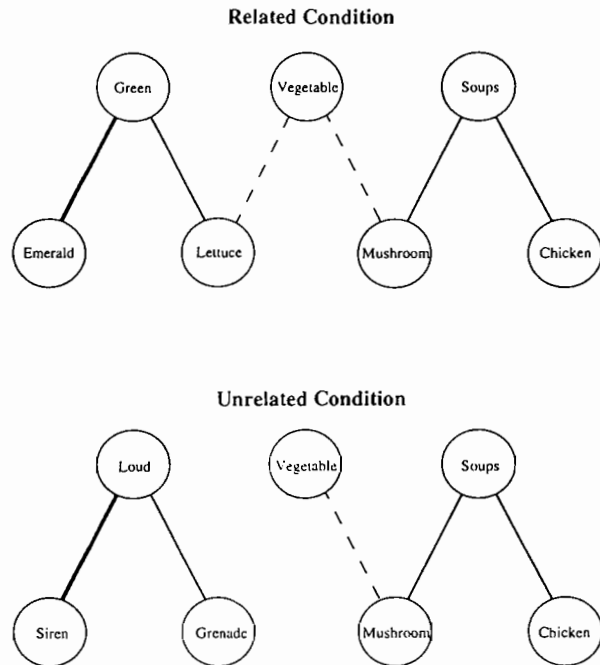


Figure 4. Design of Experiment 2. Solid lines indicate studied category–exemplar pairs; heavy lines indicate the subset of those pairs that received retrieval practice; dashed lines indicate a preexisting category–exemplar relation that was not studied in the experiment. Note that some items in the two categories are similar by virtue of the implicit category VEGETABLE instead of being members of each other’s explicit categories.

similar to the Rp– item GREEN–LETTUCE because lettuce and mushrooms both happen to be vegetables (see Figure 4). Thus, unlike in Experiment 1, we achieved the similarity between Nrp–Similar items and the Rp category by ensuring that Rp– and Nrp–Similar items were both members of a shared implicit category to which neither Rp+ nor Nrp–Dissimilar items belonged (as illustrated by the dashed lines in Figure 4). With this arrangement, participants could no longer confuse Nrp–Dissimilar items as members of the Rp category, thus eliminating reporting bias. We were then able to interpret any remaining impairment of Nrp–Similar items as the effects of inhibitory processes.

Our new Nrp–Similar items differed from the corresponding items in Experiment 1 in another important respect: There is no obvious reason why these exemplars should interfere with the retrieval practice of Rp+ items. For example, the item MUSHROOM has no a priori semantic connection to the category GREEN, so MUSHROOM should not interfere during the retrieval practice of GREEN–EMERALD. According to the inhibition hypothesis (as framed in the introduction), the Nrp–Similar item SOUPS–MUSHROOM should, therefore, remain unimpaired. If, however, retrieval practice impaired Nrp–Dissimilar items in Experiment 1 merely because those items had the misfortune of being similar to suppressed Rp– items, we should replicate that Nrp–Dissimilar impairment in the present experiment in the form of impaired recall of Nrp–Similar items. Expressed

more vividly, repeatedly retrieving GREEN–EMERALD should have the rather bizarre consequence of impairing the recall of MUSHROOM, an item that is not in any way similar to the retrieved item and that is studied and tested under a different category cue (i.e., SOUPS).

Method

The method was identical to that of Experiment 1 with the following exceptions.

Participants

The participants were 54 students from the same introductory psychology course one quarter later. None had participated in Experiment 1.

Materials and Design

Categories

We constructed three new pairs of related categories (see Appendix B). All six exemplars of each category were semantically members of only that explicit category. Some of the items in each pair of related categories, however, were similar because they shared an implicit category (a category not explicitly presented in the experiment to which those items and only those items belonged). Each category in a pair contained three items that were members of the implicit category and three items that were not. For example, GREEN and SOUPS were related categories, and VEGETABLE was their shared implicit category. GREEN included items such as EMERALD that were not members of the shared implicit category and items such as LETTUCE that were; SOUPS included items such as CHICKEN that were not members of the shared implicit category and items such as MUSHROOM that were. These exemplars were constructed according to all of the same constraints imposed on the exemplars used in Experiment 1.

Learning Booklets

We created the learning booklets as we did in Experiment 1. Again, all categories appeared equally often in the related and unrelated conditions; however, we did not need to counterbalance the membership of Rp– and Nrp–Similar items across categories (e.g., LETTUCE could never appear as a soup).

Retrieval-Practice Booklets

We constructed retrieval-practice booklets as in Experiment 1. Each booklet contained practice trials for half of the categories (one related and one unrelated). Across booklets, all categories appeared equally often as Rp and Nrp categories. Three different random practice orders were created for each set of practiced items. The learning–practice combinations were thus analogous to those shown in Table 1 for Experiment 1, but with the new Experiment 2 words.

Recall Booklets

We constructed recall booklets as in Experiment 1, except that (a) each booklet consisted of five critical pages because *mountains*—a filler category for all participants—appeared on the first page and (b) three recall orders were constructed for each learning booklet. As before, we equated the average position of related and unrelated and practiced and unpracticed categories.

Implicit Category Recall Test

In addition to cuing participants with the explicitly presented categories, we also tested their ability to recall studied exemplars when cued with the shared implicit categories to which similar items (whether in the Rp- or Nrp-Similar condition) belonged. For instance, when participants studied either or both of the categories GREEN and SOUPS, we tested their memory for VEGETABLES—the implicit category shared by the similar items of GREEN and SOUPS—by providing VEGETABLES as a cue and giving them 30 s to recall any previously studied items that fit this category. This test took place after the tests for explicit categories. It began with an instruction page informing participants that many previously studied exemplars could fit categories other than those under which they were studied and that we wanted them to try to recall as many exemplars as they could that fit each novel category. We tested each participant on three such implicit categories (one for the related categories and one each for the unrelated categories), one category per page.

Results and Discussion

All analyses were performed with learning booklet and items practiced as between-subjects counterbalancing conditions.

Retrieval Practice

As in Experiment 1, participants practiced six experimental items three times each (a total of 18 relevant practice trials). Again, the retrieval-practice success rate for items in the unrelated condition (78%) did not differ significantly from that for items in the related condition (72%), $F(1, 48) = 1.84, n.s.$

Final Recall

Note, for purposes of comparing recall percentages, that in this design (as in Experiment 1) the Rp+ and Nrp-Dissimilar conditions contained the same items and the Rp- and Nrp-Similar conditions contained the same items (see Table 2).

We report the results of planned comparisons. As before, we first verified that the retrieval-practice procedure had the expected positive and negative effects on Rp+ and Rp- items, respectively, by restricting comparisons to the unrelated condition. As expected, retrieval practice significantly facilitated the later recall of Rp+ items ($M = 72%$) relative to Nrp-Dissimilar baseline items ($M = 43%$), $F(1, 48) = 43.18, p < .001$, whereas it impaired recall of Rp- items ($M = 42%$) relative to Nrp-Similar baseline items ($M = 52%$), $F(1, 48) = 4.74, p < .05$. These findings replicate the facilitation and impairment observed in the unrelated condition of Experiment 1, as well as in prior studies (Anderson et al., 1994).

Support for Inhibitory Mechanisms

As before, the crucial question concerned the impact of retrieval practice on Nrp-Similar items in the related condition. Planned comparisons revealed that retrieval practice impaired the recall of Nrp-Similar items in the related condition relative to performance on those same items in the unrelated condition (see the boxed items in the middle of Table 2), $F(1, 48) = 12.05, p < .01$. That is, retrieval practice on Rp+ items (e.g., GREEN-EMERALD) impaired both Rp- (e.g., GREEN-LETTUCE) and Nrp-Similar items (e.g., SOUPS-MUSHROOM)

even though participants encoded and retrieved Nrp-Similar items under a different category cue and the affected items bore no relationship to the Rp category label or to the Rp+ items themselves. Thus, we replicated the cross-category impairment predicted by inhibitory models of memory retrieval.

If the cross-category inhibition evident in the impairment of Nrp-Similar items truly reflects the consequences of inhibitory mechanisms, impairment of Nrp-Similar items should generalize to novel retrieval cues. Data from the implicit category test administered at the end of the final recall phase support this prediction. When we reanalyzed the final recall data to include additional items recalled during this implicit category test, we obtained the same two effects as reported previously: Performance on Nrp-Similar items in the unrelated condition ($M = 54.3%$) was still superior to performance on Rp- items in the unrelated condition ($M = 45.1%$), $F(1, 48) = 3.86, p = .055$ (thus showing within-category inhibition), and to performance on Nrp-Similar items in the related condition ($M = 39.5%$), $F(1, 48) = 11.57, p < .01$ (thus showing cross-category inhibition). Overall, cuing with the implicit categories helped participants very little. Thus, it appears that participants' impaired ability to recall Nrp-Similar and Rp- items does not depend at all on the category cue by which those items are tested. We return to this observation in the section titled "Implications."

Evidence Against Alternative Accounts

The present findings rule out two alternative accounts of the cross-category inhibition observed in Experiment 1: reporting bias and cue bias. Consider first the reporting-bias argument, according to which cross-category inhibition derives from participants' reluctance to report correctly retrieved Nrp-Similar items in the related condition. Although reporting bias may explain the cross-category impairment observed in Experiment 1, the materials of the present study render this account implausible. For instance, it seems unlikely that participants would withhold the Nrp-Similar item MUSHROOM because they were uncertain about whether it had appeared with GREEN. Nothing in the pattern of intrusions suggests that participants experienced confusion about the membership of Nrp-Similar items or that they were excessively conservative in their reporting criterion; of a total of 702 responses across 54 participants, there were 70 intrusions (41 vs. 29 items in related vs. unrelated conditions, respectively), with 30 synonym errors (e.g., recalling "pants" instead of "slacks"), 39 intrusions of extra-experimental exemplars, and only 1 case in which a participant recalled an exemplar to the wrong category cue. As such, even if participants experienced confusion (which is highly unlikely), it seems implausible that this confusion led them to withhold the Nrp-Similar items. Furthermore, the failure of the implicit category cues provided at the end of experiment to significantly improve participants' recall corroborates these assumptions: If participants had recalled and withheld items during tests of the explicit categories, the implicit category tests should have eliminated such reluctance because they directed participants to ignore the explicitly studied category memberships. Thus, the impairment of Nrp-Similar items appears to reflect a genuine deficit in accessibility of those items.

The present Nrp-Similar impairment rules out yet another account of the cross-category inhibition observed in Experiment 1: cue bias. It has been suggested that Nrp-Similar items in Experiment 1 may have suffered because the functional representation of the Nrp category cue was changed so that it was a less potent cue for Nrp-Similar items in the related condition. Suppose that, during the final recall of the related Nrp category (e.g., FOODS) of Experiment 1, Rp- exemplars (e.g., RED-TOMATO) intruded. Such intrusions might be especially likely if Rp- items were actually primed (instead of inhibited) by the earlier retrieval practice of Rp+ items (e.g., RED-BLOOD) but were kept from being recalled within their own category by occlusion mechanisms such as those reviewed in the introduction. If Rp- exemplars intruded, they might have altered the functional representation of the Nrp category cue toward a meaning consistent with Rp- exemplars (e.g., toward fruity foods rather than starchy foods), causing the impaired recall of Nrp-Similar (e.g., RADISH) and Nrp-Dissimilar (e.g., BREAD) items.

Although the cue-bias account might explain the results of Experiment 1, it cannot handle the present cross-category inhibition. For the cue-bias account to work, one must assume that the Nrp category provides a sufficiently good retrieval cue so that Rp- exemplars will intrude during Nrp recall. This condition was not true in Experiment 2, in which Rp- exemplars (e.g., LETTUCE) were not related to Nrp categories (e.g., SOUPS) in any direct way. Because Rp- exemplars are unlikely to intrude during Nrp recall, it seems unlikely that the Nrp-Similar impairment would be caused by this factor. Furthermore, even if Rp- items had intruded, they would bias the Nrp category cue in favor of Nrp-Similar items—by virtue of the implicit category shared by Rp- and Nrp-Similar items—and away from Nrp-Dissimilar items. As a result, Nrp-Similar items would be facilitated and Nrp-Dissimilar items impaired, quite the opposite of what was observed. Thus, the present Nrp-Similar impairment rules out a cue-bias account of cross-category inhibition.

Unanticipated Results of Experiment 1 Replicated

The decreased accessibility of Nrp-Similar items in the present experiment conceptually replicates the unexpected deficit in the recall of Nrp-Dissimilar items in Experiment 1. As with the Nrp-Dissimilar items in that earlier study, the present Nrp-Similar items had no direct similarity whatsoever to the practiced items themselves and had no a priori semantic association to the Rp category; however, they were directly similar to the Rp- items, which we assume were vulnerable to inhibition during the retrieval practice of Rp+ items.

This conceptual replication is surprising in light of changes in materials that might have worked against a replication; in particular, introducing a shared implicit category (e.g., VEGETABLES) as the basis of the similarity between Rp- and Nrp-Similar items might have obscured the crucial similarity relationships that were more obvious in Experiment 1. In Experiment 1, the similarities between the Rp- and Nrp-Dissimilar items were highlighted by the study conditions. For example, the presence of FOOD as a study category should have increased the salience of that concept, making it more likely to be spontaneously encoded as part of the representation for the

Rp- item RED-TOMATO. The presence of the shared category FOOD in the representation of TOMATO would make that item similar to the Nrp-Dissimilar item FOOD-CRACKERS. In contrast, the shared category in Experiment 2 (e.g., VEGETABLES) remained implicit, decreasing the likelihood that participants would spontaneously encode it in their representations of both Rp- and Nrp-Similar items. If participants had failed to encode this implicit category into their functional representations, impairment of Nrp-Similar items might not have been observed.

This replication of the Nrp-Dissimilar impairment shows that retrieval impairs not only the items that directly interfere with access to the retrieval-practice target items but also the items that are merely similar to those interfering competitors. The conspicuous lack of impairment for Nrp-Dissimilar items in the present experiment ($M = 43\%$ for both related and unrelated conditions) reinforces this conclusion; unlike those items bearing the same name in Experiment 1, Nrp-Dissimilar items in the present experiment bore no resemblance to Rp- items and thus should not have been impaired according to this new similarity hypothesis. Thus, inhibition appears to selectively extend to items directly similar to suppressed exemplars.

EXPERIMENT 3: NECESSARY CONDITIONS FOR INHIBITION

The results of Experiments 1 and 2 favor an inhibitory theory of retrieval-induced forgetting over all of the noninhibitory alternatives reviewed in the introduction. Retrieval practice impairs the recall of competing items regardless of whether they share a common retrieval cue with those practiced items, and this impaired performance reflects a genuine reduction in the accessibility of those items. Yet, the impairment of Nrp-Dissimilar items in Experiment 1 and of Nrp-Similar items in Experiment 2 remains a mystery, even from the standpoint of the inhibition model. The impairment of these items clearly seems to derive from both the retrieval practice of Rp+ items and the similarity of those Nrp items to Rp- items directly vulnerable to inhibition. Still, there is no obvious way in which these items might cause interference during retrieval practice, so it is unclear why they should be impaired.

To further explore the basic properties of this puzzling impairment, we decided to test our assumptions regarding the necessary conditions for its occurrence: retrieval practice and similarity to other items undergoing inhibition. If our characterization of this impairment is correct, eliminating either of these conditions should eliminate the impairment. This was the strategy of Experiments 3a and 3b.

Experiment 3a: Retrieval Practice as a Necessary Condition

The purpose of this experiment was to determine whether performing retrieval practice on Rp+ items was truly necessary for the recall impairment of Nrp-Similar items in Experiment 2. Perhaps the mere simultaneous presence of the two semantically related categories in memory might have caused impairment of Nrp-Similar items, independently of whether participants practiced Rp+ items. The simultaneous presence of the

two related categories might have impaired the recall of Nrp-Similar items in two ways. First, it is possible that, during the study phase, participants noticed the implicit semantic relationship between Rp- and Nrp-Similar exemplars in the related condition; attending to this implicit category might have distracted participants from encoding the relationship of those items to their respective explicit categories, causing the later recall of Nrp-Similar exemplars to be impaired. If participants were less likely to be distracted by the implicit category in the unrelated condition, this encoding deficit might explain the apparent inhibition of Nrp-Similar items in the related condition.

A second way that studying the two related categories might have caused impairment of Nrp-Similar items would be if (a) participants used the implicit category (e.g., VEGETABLES) on the final test when they searched for Nrp-Similar items in memory and (b) the presence of Rp- items in the related condition increased the number of memory targets associated with that implicit category (e.g., it increased the number of studied vegetables in memory). These conditions might have reduced the recall of Nrp-Similar items in the related condition by increasing the effects of cue overload (Watkins, 1975, 1978) in the related condition relative to the unrelated condition, in which Rp- items were not members of the implicit category. A similar argument can be made with respect to Experiment 1 (in which, for example, there were nine red things in memory when RED was in the related condition but only six when RED was in the unrelated condition). Because these experiments confounded our manipulation of category relatedness with the number of items that could be associated with a real or implicit retrieval cue, we cannot confidently ascribe the impairment of Nrp-Similar items to retrieval practice of Rp+ items.

To evaluate this possibility, we replicated Experiment 2 with one minor change. Instead of performing retrieval practice on the critical Rp+ items, as in Experiment 2, participants practiced only filler category items. If the mere presence of the related category in memory impairs the recall of Nrp-Similar items, Nrp-Similar recall should be lower in the related condition than in the unrelated condition, even though participants did not practice items in the related category. If, however, retrieval practice is necessary to impair Nrp-Similar items, failure to practice Rp+ items from a related category should eliminate the impairment observed in Experiment 2.

Method

Participants

The participants were 27 students from the same introductory course as in Experiment 2. None had participated in any other inhibition experiment.

Materials and Design

Categories, learning booklets, and recall booklets. Retrieval-practice booklets contained 51 pages for practicing items from filler categories only: only one version was created. The rules of construction of the retrieval-practice booklets for Experiment 2 were loosely followed. Because there were only four filler categories and participants practiced no item more than three times, participants practiced more than half of the items from each of the filler categories.

Design. Unlike Experiments 1 and 2, this experiment used a 2×2 within-subjects design. Participants studied related and unrelated categories but did not practice them; thus, there were only Nrp-Dissimilar and Nrp-Similar items.

Results and Discussion

In this experiment, there were no Rp+ or Rp- items because participants practiced only filler categories. Participants therefore provided twice as much data for the Nrp-Dissimilar and the Nrp-Similar conditions because, for each participant, two categories contributed to each cell (see Table 2). All analyses were performed with learning booklet as a between-subjects counterbalancing condition.

Retrieval-Practice and Intrusion Data

Because participants practiced only filler items, we did not score retrieval-practice performance in the present experiment. The intrusion rate on the final recall test was comparable to that of the previous experiments, with 41 intrusions out of a total of 323 responses (an intrusion rate of 12.7%).

Final Recall

The experiment revealed no evidence suggesting that the mere presence of the two related categories in memory affected overall recall in general, and there was no main effect of relatedness ($M_s = 42\%$ for unrelated and 46% for related), $F(1, 24) = 1.53$, *ns*. Moreover, no evidence suggested that the mere presence of the two related categories in memory affected recall for Nrp-Similar items specifically, $F(1, 24) < 1$. The only statistically significant effect in the present experiment was an overall difference between the recall of Nrp-Similar and Nrp-Dissimilar items, $F(1, 24) = 5.57$, $p < .05$, indicating that participants recalled Nrp-Similar items more easily ($M = 48\%$) than Nrp-Dissimilar items ($M = 40\%$). This finding replicates the analogous finding in Experiment 2 (we discuss this pattern in the *Results and Discussion* section of Experiment 3b). Finally, there was no interaction between relatedness and item type, $F(1, 24) = 2.16$, *ns*. Thus, the present data support neither the encoding deficit nor the cue-overload interpretations of Nrp-Similar impairment.

Experiment 3a was designed as a control condition for Experiment 2; we ran it at the same time as Experiment 2 and used the same materials. We are therefore justified in comparing the present Nrp-Similar recall performance with that obtained in Experiment 2. Retrieval practice of Rp+ items appeared to be a necessary condition for impairing the recall of Nrp-Similar items in the design of Experiment 2. The interaction of experiment and category relatedness with respect to Nrp-Similar recall performance was significant (compare the boxed numbers in the middle and bottom of Table 2), $F(1, 60) = 3.98$, $p < .05$, providing further confirmation that Nrp-Similar items were reliably impaired in Experiment 2 (in which retrieval practice was present) but not in Experiment 3a, in which there was no retrieval practice on related exemplars.

Experiment 3b: Similarity as a Necessary Condition

The purpose of Experiment 3b was to determine whether similarity to inhibited Rp- items was a necessary condition for the impaired recall of Nrp-Similar items in Experiment 2. After critical reexamination of the materials from Experiment 2, we discovered a few cases in which one might argue that Rp+ items were directly similar to Nrp-Similar exemplars. One of these embarrassing cases comes from the example that we have been using to illustrate the materials from Experiment 2: CHICKEN (a member of SOUPS) is similar to LETTUCE (a member of GREEN) in that they are both foods. Thus, when participants did retrieval practice on SOUPS-CHICKEN, that Rp+ item was directly similar to the Nrp-Similar item GREEN-LETTUCE. In this instance, impairment of GREEN-LETTUCE, although still interesting as an example of cross-category inhibition, would not reflect impairment resulting from the mere similarity of the Nrp-Similar item to Rp- exemplars. Most of our materials do not appear subject to this criticism; nonetheless, hidden similarities between Rp+ and Nrp-Similar items may have enabled retrieval practice to impair Nrp-Similar items, creating the illusion of impairment deriving from mere similarity to inhibited Rp- items.

To evaluate this hypothesis regarding the impaired recall of Nrp-Similar items, we replicated Experiment 2 with one change in the materials. For each pair of categories (e.g., GREEN and SOUPS), we constructed a set of items that could replace the present Rp- exemplars. The new replacement exemplars were dissimilar items; that is, they belonged to one studied category but did not share an implicit category with any members of the previously related category. For example, for the participants who practiced GREEN, we replaced LETTUCE (an Rp- item) with FROG. As we illustrate in Figure 5, replacing LETTUCE with FROG "lesions" the similarity between Rp- items and Nrp-Similar items (e.g., SOUPS-MUSHROOM). Thus, if Rp+ and Nrp-Similar exemplars share hidden similarities, retrieval practice of GREEN-EMERALD should still impair the later recall of SOUPS-MUSHROOM, even though we replaced LETTUCE with FROG. However, if Nrp-Similar items must resemble Rp- exemplars for impairment to occur, practicing Rp+ items should not impair the recall of Nrp-Similar items.

Method

The method was identical to that of Experiment 2 with the following exceptions.

Participants

The participants were 48 students from a later introductory psychology course; none had participated in any of the other experiments.

Design and Materials

Design. This experiment also used a 2×4 within-subjects design, but the independent variables differed from those in Experiments 1 and 2. First, we manipulated the practice status of the category pair. Participants always learned both related categories of a given pair, however, in the *practiced-pair* condition they practiced one category of the pair, whereas in the *unpracticed-pair* condition they practiced neither category of the pair. For example, participants might have studied the cate-

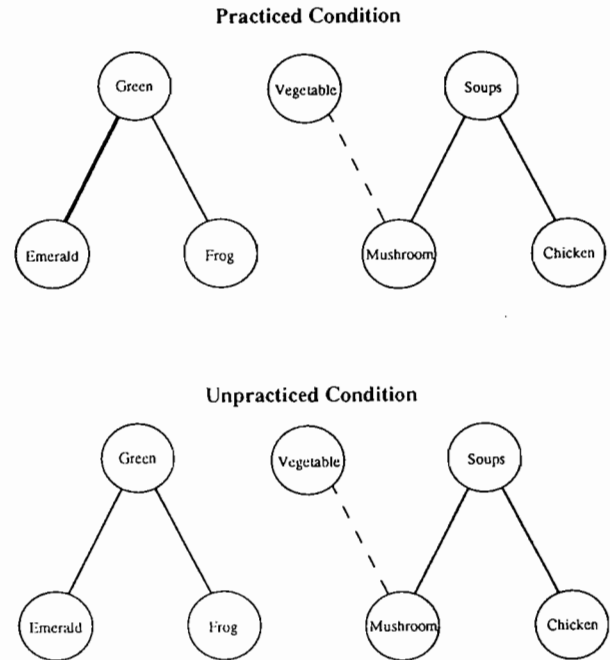


Figure 5. Design of Experiment 3b. Solid lines indicate studied category-exemplar pairs; heavy lines indicate the subset of those pairs that received retrieval practice; dashed lines indicate a preexisting category-exemplar relation that was not studied in the experiment. Note that the practiced category no longer is related to the implicit category VEGETABLE.

gory pairs GREEN-SOUPS and SHARP-LOUD. If participants practiced only the GREEN category, then GREEN and SOUPS would constitute the practiced-pair condition, and SHARP and LOUD would constitute the unpracticed-pair condition. Second, we manipulated the type of items from which we composed categories. There were altered categories composed of dissimilar and replacement items and non-altered categories composed of dissimilar and similar items. Participants practiced only the dissimilar items from one altered category. For instance, in the preceding example, if GREEN and SHARP were altered categories and SOUPS and LOUD were nonaltered categories, then participants would practice only the dissimilar items—GREEN-EMERALD, GREEN-DOLLAR, and GREEN-LAWN—from only one of the altered categories.

The rationale for the present change in design requires clarification. In previous designs, we used the unrelated condition to provide a baseline against which to measure the within-category impairment of Rp- items and the cross-category impairment of Nrp-Similar items in the related condition. In the present experiment, the unpracticed-pair condition provided the necessary baselines. After replacing similar Rp- items such as LETTUCE with replacement items such as FROG, we needed to determine two things: (a) whether retrieval practice in fact impaired the new replacement items (which could be determined by comparing recall for replacement items in the practiced- and unpracticed-pair conditions) and (b) whether retrieval practice impaired Nrp-Similar items when we lesioned their similarity to Rp- items (which could be determined by comparing recall for Nrp-Similar items in the practiced- and unpracticed-pair conditions). Thus, by providing participants with both categories of a pair (one altered and one not) and then having them practice none of those items in the unpracticed-pair condition, we could obtain the necessary baselines for both the new replacement items and Nrp-Similar exemplars (see Figure 5).

Items. We found three replacement items for each of the six cate-

ries (see Appendix B for a list). Replacement items met the following criteria: (a) They were members of the category under which they were studied; (b) they were not members of the shared implicit category and, therefore, were not similar to Nrp-Similar items; and (c) the three new items did not, themselves, form some kind of subcategory. Replacement items were also subject to the same first-two-letter constraints as the other materials.

Learning booklets. Participants again learned eight (four experimental and four filler) categories. This time, however, they learned two pairs of related categories. One member of each related pair contained the same similar and dissimilar items as in the previous experiments; the other member of the related pair contained dissimilar and replacement items. For example, in one learning booklet, SOUPS retained all of its original exemplars (e.g., CHICKEN and MUSHROOM), whereas its related category, GREEN, retained its dissimilar items (e.g., EMERALD) but had its similar items (e.g., LETTUCE) replaced with new items such as FROG.

Retrieval-practice booklets. The present experiment differed from previous ones in that participants performed retrieval practice on dissimilar items from only one experimental category instead of two. The practiced category was one of the two categories that had dissimilar and replacement items in the learning booklets. For example, if the learning booklet contained GREEN, SOUPS, LOUD, and SHARP, and the GREEN and LOUD categories had items replaced, then participants practiced either GREEN or LOUD. Practice booklets again contained 51 practice trials; we constructed the booklets by making two adjustments to those of Experiment 2. First, we replaced one of the experimental categories with a filler category so that participants practiced only one experimental category. Second, we used only one practice order for each practiced category.

Recall booklets. Each booklet began with the category *mountains*. We inserted the one practiced category in all four possible positions with the following constraints: Related categories never appeared adjacently and, across all recall orders, the average position of the practiced and unpracticed pairs of categories was equated. As a result of those constraints, the average position of categories containing similar and replacement items was equated.

Results and Discussion

The results appear in Table 3. Note, for purposes of comparing recall percentages, that the two columns with "dissimilar" in the heading contain the same items; the other two columns do not contain the same items (replacement and similar items

were different). All analyses were performed with learning booklet as a between-subjects counterbalancing condition.

Retrieval Practice

Participants practiced three critical items three times each (a total of nine practice trials) and had a success rate of 79%. This percentage reflects the real data of 42 participants and the estimated data for 6 participants whose retrieval-practice booklets were missing. We estimated missing practice data by inserting the Rp+ final recall score for that participant into the analysis, because final recall and retrieval-practice success are highly correlated.

Final Recall

As in Experiments 1 and 2, retrieval practice facilitated the final recall of practiced items ($M = 72\%$) relative to their corresponding baseline items (for dissimilar-unpracticed items from a category with replacement exemplars, $M = 38\%$), $F(1, 42) = 50.08$, $p < .001$, whereas it impaired the recall of within-category Rp- competitors (for replacement items from the practiced-pair condition, $M = 35\%$) relative to their corresponding baseline items (for replacement items from the unpracticed-pair condition, $M = 46\%$), $F(1, 42) = 6.95$, $p < .01$. Thus, retrieval practice impaired our new replacement Rp- items (e.g., GREEN-FROG) just as effectively as it impaired previous Rp- items. Table 3 displays these findings.

Nrp-Similar impairment eliminated. The crucial question in the present experiment concerned the fate of Nrp-Similar items as a consequence of practicing Rp+ exemplars. A planned comparison demonstrated that retrieval practice did not significantly impair Nrp-Similar items in the practiced-pair condition ($M = 47\%$) relative to those same items in the unpracticed condition ($M = 49\%$), $F(1, 42) < 1$. Thus, lesioning the similarity between Rp- items (e.g., GREEN-LETTUCE) and Nrp-Similar items (e.g., SOUPS-MUSHROOM) by replacing Rp- items with nonsimilar replacement items (e.g., GREEN-FROG) eliminated the impairment of Nrp-Similar exemplars, indicating that their impairment in Experiment 2 was (a) not produced

Table 3
Experiment 3b: Mean Percentage of Words Recalled as a Function of Practice Status of the Category Pair and the Item Types From Which Categories Were Composed

Practice status of category pair	Item type			
	Categories altered with replacement items (GREEN)		Categories not altered with replacement items (SOUPS)	
	Dissimilar (EMERALD)	Replacement (FROG)	Similar (MUSHROOM)	Dissimilar (CHICKEN)
Unpracticed pair	38	46	49	43
Practiced pair	72+	35-	47	42

Note. $N = 48$. The box shows the comparison for detecting cross-category inhibition. A+ designates the only items in the experiment to receive retrieval practice (i.e., Rp+ items); A- designates the only items in the experiment whose category mates received retrieval practice (i.e., Rp- items).

by any hidden similarity to Rp+ items and (b) critically dependent on their similarity to inhibited Rp- items.

Role of nominal versus functional similarity in inhibition. The present findings clearly demonstrate that similarity considerations are important in predicting when this unusual variety of cross-category inhibition will occur. But in what way should these cross-category items be similar to Rp- items? Throughout this article, we have operationalized similarity simply by membership in a common category. Careful thought will reveal, however, that the category membership criterion for similarity is not sufficient by itself to accurately predict when impairment will and will not occur. Consider the following argument about why retrieval practice of GREEN-EMERALD should impair SOUPS-CLAM: Emeralds are similar to pearls (a nonstudied item) in that they are both members of the category *precious items* (a nonstudied category); thus, if category membership is sufficient, retrieval practice on GREEN-EMERALD should inhibit *pearl*. Pearls are similar to clams in that they both are members of the category *found-in-the-sea* (a nonstudied category). Thus, because emeralds are similar to pearls and because pearls are similar to clams, CLAM should be inhibited by retrieval practice of GREEN-EMERALD. Yet CLAM is not inhibited. Why not?

One reason why CLAM should not be impaired derives from the distinction between nominal similarity (i.e., similarity between items as construed by the experimenters) and functional similarity (i.e., similarity as represented by the participants). Experimenters can fabricate a basis on which nearly any two verbal items might be similar (e.g., the category WAR THINGS renders GRENADE and MUSHROOM similar if one thinks of "mushroom clouds" instead of the edible referent). From the participants' standpoint, however, such fabrications are irrelevant unless they enter into their functional representations of the items (i.e., unless the participants encode them). In the pearl-clam example noted previously, participants might not spontaneously encode FOUND-IN-THE-SEA given the stimulus SOUPS-CLAM unless there were enough *sea* items on the study list to make that property salient. In contrast, participants in our study were very likely to represent the critical similarity between Nrp-Similar and Rp- items because (a) the implicit categories on which those items are similar were constructed to be fairly standard categories for those items (e.g., VEGETABLES, CLOTHING, and WEAPONS) and (b) six members of the implicit category appeared on the study list across two explicit categories. Both of these factors should have increased participants' likelihood of encoding that particular similarity between Nrp-Similar and Rp- items, an assumption compatible with the consistent advantage of Nrp-Similar items over Nrp-Dissimilar items in the unrelated condition, the latter of which had no implicit category by which their recall may have been facilitated. (This latter interpretation is complicated, however, by potential differences in item difficulty across these two sets.) Thus, the likelihood that participants functionally represent a property that might form a basis for similarity should determine whether Nrp items will be impaired.

Competitor inhibition as a prerequisite of Nrp-Similar impairment. In the previous section, we highlighted that, for Nrp-Similar impairment to occur, participants must have a functional representation of the Nrp-Similar item that is similar

to an item whose competitors are practiced. The findings of the present experiment underscore that this condition, by itself, is not sufficient to cause impairment. Impairment of an Nrp-Similar item depends on both its similarity to an item whose competitors are practiced and the successful inhibition of that similar item.

Consider the fact that impairment of the item SOUPS MUSHROOM was eliminated when GREEN LETTUCE was replaced by GREEN FROG. Even though GREEN LETTUCE did not appear on the study list, we assumed that MUSHROOM and LETTUCE were still similar from the participants' perspective in that they are both considered vegetables. Furthermore, we assumed that participants still knew that lettuce is green. Why, then, was *mushroom* not impaired? Presumably, MUSHROOM was not impaired because LETTUCE did not appear on the study list and was thus not sufficiently active to cause interference during practice of GREEN EMERALD. As such, LETTUCE would not have been inhibited if the function of inhibition is to overcome interference. Thus, the present experiment shows that, to be impaired, Nrp-Similar items must be similar to a competitor of the practiced item that is actually inhibited.

EXPERIMENT 4: ENDURANCE OF THE INHIBITION

Experiments 1-3 strongly suggest that the retrieval-induced forgetting observed in the retrieval-practice paradigm results from an inhibitory process that assists in the resolution of retrieval competition during the retrieval-practice phase. If this interpretation of retrieval-induced forgetting is correct, then the fact that such impairment is still measurable after a 20-min retention interval becomes especially important because the temporal durability of inhibitory processes is normally thought to be far briefer. Evidence that the consequences of inhibitory processes could endure for as long as 20 min would be highly unusual and perhaps informative with respect to both the specific nature of the impairment and the characteristics of cognitive inhibition more generally.

Unfortunately, the present experiments do not allow us to reach conclusions about the durability of retrieval-induced forgetting because all of the impairment could have arisen during the final test. For example, cuing participants to retrieve the GREEN category before cuing them to retrieve SOUPS may have enabled their retrieval of GREEN-EMERALD on the prior test page to temporarily inhibit SOUPS-MUSHROOM. Although we did not always test the practiced category first, Nrp-Similar impairment may have been large enough in those cases to offset any lack of such impairment when we tested practiced categories last. Such cross-category inhibition, although interesting from the standpoint of evaluating noninhibitory theories of interference, would not reflect a durable form of forgetting.

Some evidence already exists, however, that impairment in the retrieval-practice paradigm does in fact endure throughout the 20-min retention interval. In two studies, Anderson et al. (1994) replaced the category-cued free recall test with a category-plus-stem cued recall procedure (e.g., FRUIT O _____) precisely so that they could manipulate the order in which participants retrieved the Rp+ and Rp- exemplars in a given category. Anderson et al. found significant retrieval-induced forgetting

with this new testing procedure even when they tested Rp- items before Rp+ items in the final testing booklet. Because testing Rp- exemplars before Rp+ exemplars eliminated any within-category impairment of Rp- exemplars that might have occurred during the final test phase, Anderson et al. concluded that the negative effects of retrieval practice on Rp- items must have endured throughout the 20-min retention interval. Thus, it is possible that the impairment of Nrp-Similar items in the present experiments also reflects the enduring consequences of retrieval practice.

One might object that the Anderson et al. (1994) findings do not provide clear evidence for durable inhibition because the letter cues with which participants' final recall was directed may not have eliminated covert retrievals of competing Rp+ items. Any such covert retrievals that occurred before recall of the Rp- item may well have impaired Rp- recall at final test. The present paradigm affords a simple way of testing whether retrieval-induced forgetting endures throughout the 20-min retention interval that is not subject to this criticism. Rather than changing the final testing procedure from category-cued free recall to category-plus-stem cued recall as in the Anderson et al. study, we simply replicated our present Experiment 2, manipulating the order in which we tested the Nrp and Rp categories in the final retrieval booklet. If Nrp-Similar impairment arises entirely during the final test, we should find this impairment only when we test Rp categories before Nrp categories. If, however, Nrp-Similar impairment endures from the practice phase 20 min earlier, then impairment should occur regardless of whether we test the Nrp category before or after the related Rp category. That is, participants should recall Nrp-Similar items in the related category condition less well than the corresponding Nrp-Similar items in the unrelated condition, regardless of the order in which we test practiced and unpracticed categories.

Method

The method was identical to that of Experiment 2 with the following exceptions.

Participants

The participants were 72 students from a later introductory psychology course; none had participated in any of the other experiments.

Design

We added a third, between-subjects variable to make this a $2 \times 2 \times 4$ design. Participants were randomly assigned to either the *practiced first* condition, in which we tested the practiced categories before the unpracticed categories, or to the *unpracticed first* condition, in which we did the reverse.

Materials

Recall Booklets

For each set of learned items, we constructed four different recall booklets. Each booklet began with MOUNTAINS. In two of the booklets, the practiced categories were recalled before the unpracticed categories; in the other two, the unpracticed categories were recalled before the practiced categories. Related and unrelated categories always alternated (half of the booklets started with a related category and half with an

unrelated category) so that participants never recalled a category adjacently to its related category and the average recall position for related and unrelated categories was equated. We ensured that every category used in the experiment appeared in every position in the testing booklet, once in each condition.

Results and Discussion

All analyses were performed with learning booklet and items practiced as between-subjects counterbalancing conditions.

Retrieval-Practice and Intrusion Data

As in Experiment 2, participants practiced six experimental items three times each for a total of 18 relevant retrieval-practice trials. The success rates for items from unrelated ($M = 78\%$) and related ($M = 73\%$) categories were not significantly different, $F(1, 58) = 1.19, ns$. The intrusion rate on the final recall test was comparable to that observed in the previous studies, with 106 intrusions out of 999 responses (an intrusion rate of 10.6%).

Final Recall

Note, for purposes of comparing recall percentages, that in this design the Rp+ and Nrp-Dissimilar conditions contained the same items, as did the Rp- and Nrp-Similar conditions (see Table 4).

Again, as in Experiments 1, 2, and 3b, retrieval practice facilitated the recall of Rp+ items ($M = 70\%$) in the unrelated condition relative to the recall of corresponding Nrp-Dissimilar items ($M = 47\%$), $F(1, 60) = 43.65, p < .001$, whereas it impaired the recall of Rp- items ($M = 43\%$) relative to their corresponding Nrp-Similar baseline items ($M = 57\%$), $F(1, 60) = 12.36, p = .001$. Neither of these effects interacted with the testing-order manipulation, $F(1, 60) < 1$ in both cases. Overall recall did not differ as a function of the testing order of categories ($M = 52\%$ for both orders), $F(1, 60) < 1$.

As in Experiment 2, retrieval practice on Rp+ items impaired the recall of Nrp-Similar items in the related condition ($M = 47\%$) relative to the recall observed for the corresponding items in the unrelated condition ($M = 57\%$), $F(1, 60) = 7.53, p < .01$, further replicating the unusual Nrp-Similar impairment. The crucial question, however, concerns whether this Nrp-Similar impairment differed as a function of the order in which we tested Rp and Nrp categories in the final test booklet. The significant impairment of Nrp-Similar items did not interact with testing order, $F(1, 60) < 1$, demonstrating that retrieval practice impaired Nrp-Similar items even when we tested Nrp categories before Rp categories (compare the two boxes in Table 4). Thus, the prior output of Rp+ items at test does not cause the impairment of Nrp-Similar items; rather, the impairment reflects inhibition of those items that endured from the retrieval-practice phase.

SUMMARY OF EXPERIMENTS

The present experiments sought evidence that retrieval-induced forgetting in episodic memory arises from the active inhibition of the impaired items during the retrieval of target ma-

Table 4
Experiment 4: Mean Percentage of Words Recalled as a Function of Order of Category Recall, Category Relatedness, and the Retrieval-Practice Status of an Item

Order of recall and relatedness	Retrieval-practice status			
	Practiced category (GREEN)		Unpracticed category (SOUPS)	
	Rp+ (EMERALD)	Rp- (LETTUCE)	NrpS (MUSHROOM)	NrpD (CHICKEN)
Practiced first				
Unrelated	71	43	57	47
Related	64	39	49	44
Unpracticed first				
Unrelated	69	43	57	47
Related	67	41	45	44

Note. $N = 72$. The comparisons for detecting cross-category inhibition are shown in boxes. Unrelated categories were learned without their paired category in the learning booklet; related categories were learned with their paired category in the learning booklet. Rp+ = practiced items; Rp- = unpracticed items; NrpS = similar items from unpracticed categories; NrpD = dissimilar items from unpracticed categories. Rp+ and NrpD are the same item set; Rp- and NrpS are the same item set.

terial. To assess whether inhibitory processes contribute to retrieval-induced forgetting, we devised the independent probe technique, which takes advantage of the fact that inhibitory models make the unique prediction that retrieval will impair the recall of similar competitors studied and tested under distinct cues. Across four experiments, the independent probe technique yielded three findings that place important constraints on theories of retrieval.

First, retrieving a target item from long-term memory may impair the later recall of other memory items studied and tested under cues different from those used to retrieve the target. For example, in Experiment 1, retrieval practice on an item such as RED-BLOOD (Rp-) not only impaired the recall of its unpracticed category mates like RED-TOMATO (Rp-) but also impaired the recall of items like FOOD-STRAWBERRY (Nrp-Similar), which semantically is also a red thing but one that was studied and tested with a different category cue. This cross-category inhibition was replicated with different materials in Experiments 2 and 4.

Second, the impairment may be seen not only in items that are similar to the retrieved item but also in some more distant items; we hereafter refer to this effect as *second-order inhibition*. In Experiment 1, we discovered that, in addition to the impairment of Rp- items (RED-TOMATO) and Nrp-Similar items (FOOD-STRAWBERRY) as just described, retrieval practice on Rp+ items such as RED-BLOOD also impaired Nrp-Dissimilar items such as FOOD-CRACKERS. Such an impairment was unexpected because CRACKERS is not semantically similar to BLOOD, and it was not studied as a member of the RED category; thus, because CRACKERS should have caused no interference during the practice of RED-BLOOD, it should not have been impaired. We noted, however, that although items like FOOD-CRACKERS were not similar to practiced items like RED-BLOOD (Rp+), they were systematically similar to BLOOD's primary competitors—such as RED-TOMATO (Rp)—in that both CRACKERS and TOMATO are members of the semantic category *food* (a systematic relationship arising as an accidental by-product

of our counterbalancing measures [see Appendix A: the relationship is illustrated in Figure 2]). Thus, FOOD-CRACKERS may be regarded as having a second-order similarity to RED-BLOOD in that FOOD-CRACKERS is similar to RED-TOMATO, which, in turn, is similar to RED-BLOOD (see the *Results and Discussion* section of Experiment 3b for a discussion of similarity as it is used in the present context). Therefore, the finding that performing retrieval practice on RED-BLOOD impaired FOOD-CRACKERS can be described as second-order inhibition.

The second-order inhibition effect observed in Experiment 1 was replicated in both Experiments 2 and 4 and appears to be a direct result of those items bearing similarity to other items directly vulnerable to inhibition. In Experiments 2 and 4, performing retrieval practice on GREEN-EMERALD (Rp+) impaired the later recall of SOUPS-MUSHROOM (Nrp-Similar) merely because MUSHROOM had the misfortune of being similar to EMERALD's direct competitor GREEN-LETTUCE (Rp-) in that both belonged to the shared implicit category VEGETABLES. That the impaired recall of SOUPS-MUSHROOM derives from its similarity to the inhibited item LETTUCE was supported by two additional findings: (a) Experiment 3a showed that removing the source of inhibition of GREEN-LETTUCE (i.e., eliminating retrieval practice on GREEN-EMERALD) eliminated the impaired recall of SOUPS-MUSHROOM and (b) Experiment 3b demonstrated that replacing items like LETTUCE with nonvegetable items like FROG eliminated the ability of the retrieval practice of GREEN-EMERALD to impair SOUPS-MUSHROOM, exactly what would be expected if the impaired recall of MUSHROOM depended on its similarity to LETTUCE. Thus, merely having the misfortune of being similar to an item that was directly inhibited rendered an item vulnerable to retrieval-induced forgetting. Interestingly, this second-order inhibition occurred even when we tested recall with a novel category cue in addition to the category under which the affected item was explicitly studied, suggesting a generalized inhibition.

Third, retrieval-induced forgetting appears to endure for at least 20 min. This characteristic was clearly established in Ex-

periment 4, which manipulated the testing order of practiced and unpracticed categories to evaluate the possibility that cross-category inhibition was caused during the final recall test. For example, if the practiced category GREEN were tested before the related category SOUPS, the prior recall of GREEN-EMERALD (Rp+) might temporarily suppress SOUPS-MUSHROOM (Nrp-Similar), making it appear as though the retrieval-induced forgetting had lasted for 20 min when it had not. Experiment 4 demonstrated, however, that cross-category impairment occurs regardless of the order in which the practiced and unpracticed categories are tested, supporting our belief that the inhibition endured throughout the entire 20-min retention interval (see Anderson et al., 1994, for a related finding).

Each of these findings is highly counterintuitive. The prediction of cross-category inhibition, for example, should have failed to materialize on the basis of a couple of generalizations that we regard as common wisdom among memory researchers. The encoding specificity principle (Tulving & Thomson, 1973), for instance, would lead one to expect that encoding the item STRAWBERRY under the retrieval cue FOOD would lead to a highly specific encoding that would "protect" STRAWBERRY from any harm that might be caused by the practice of RED-BLOOD; also, classical findings relating the degree of retroactive interference between two memory items to the similarity of the cues used to encode those items lead one to expect no cross-category inhibition (Osgood, 1949).² The data, however, seem unambiguous: Retrieval practice causes cross-category inhibition of surprising magnitude and tenacity. In the next section, we consider in more detail the implications of these findings for the debated role of inhibitory processes in memory retrieval. We then take a broader view of our findings, discussing their theoretical and empirical similarity to claims of cognitive inhibition outside of memory retrieval.

IMPLICATIONS

Inhibitory Processes in Memory Retrieval: A Debate Resolved?

Until now, the status of inhibitory processes in memory retrieval has been a matter of theoretical parsimony. If the impaired recall of a memory item arising from the retrieval of its competitor can be explained with only excitatory mechanisms, there seems no point in postulating additional inhibitory processes. We believe, however, that the present findings clearly favor the conclusion that retrieval uses inhibitory mechanisms in the resolution of retrieval competition. Before drawing that conclusion, we consider whether existing noninhibitory theories could be adapted to accommodate our findings. Then we consider whether inhibitory mechanisms accommodate both the cross-category and second-order inhibition effects better than noninhibitory theories. We close our discussion of this debate by clarifying the important empirical dimension along which inhibitory and noninhibitory models differ—cue-independent forgetting—and by describing conditions that may be necessary for inhibition to occur.

Noninhibitory Models Revisited

According to noninhibitory theories of interference, retrieval-induced forgetting arises from associative factors specific to the category retrieval cues used during retrieval practice. Specifically, retrieval practice on a particular cue-target pair (e.g., RED-BLOOD) should strengthen the cue-target association, occluding, diverting activation resources from, or decreasing the associative link to other competitors sharing that retrieval cue (e.g., RED-TOMATO). Because using an independent retrieval cue (e.g., FOOD) circumvents these associative dynamics, noninhibitory theories cannot accommodate cross-category inhibition in any simple way. In this section, we consider more complex interpretations of our task that might still be consistent with these noninhibitory theories.

Influence of Multiple Category Cues

One approach to explaining cross-category inhibition emphasizes additional strategies or cues that participants might have adopted as a consequence of our materials. Although we have addressed several such hypotheses in our discussions of Experiments 1, 2, 3a, and 3b, we address one final class of possibilities that is of particular interest to noninhibitory theories of interference: the influence of the availability of multiple category cues.

In Experiment 1, suppose that, in addition to encoding explicitly provided category-exemplar relationships, participants noticed our hidden between-categories similarities. For example, suppose that while studying the Nrp-Similar item FOOD-STRAWBERRY, participants noticed that STRAWBERRY could also be a RED thing, prompting them to associate the RED category with the FOOD category cue. Later, when cued with the FOOD category, any attempt to use RED as an additional category cue to recall STRAWBERRY might have backfired on participants, allowing RED-BLOOD to block (or rob resources from) the critical item and impair its recall. Because RED-BLOOD was neither studied nor practiced in the unrelated condition, similar blocking could not have occurred. Thus, the difference between Nrp-Similar items in the related and unrelated conditions might be explainable with noninhibitory interference mechanisms.

² Scholars of interference theory may sense an apparent discrepancy between the present evidence for cross-category inhibition and classic findings in the interference domain, as depicted by Osgood's transfer surface (Osgood, 1949). The present findings are indeed quite unexpected given that surface, but they do not contradict it. For instance, according to Osgood, there should be very little retroactive interference between two paired associates (S1-R1 and S2-R2) that have entirely dissimilar stimulus terms (S1 and S2), even when their respective responses (R1 and R2) are similar (e.g., RED-BLOOD and FOOD-STRAWBERRY). Our findings appear to contradict this generalization because strengthening the S1-R1 association (RED-BLOOD) impairs the ability to recall R2 (STRAWBERRY), given the stimulus S2 (FOOD), even though the RED and FOOD stimuli are dissimilar. An important difference, however, is that the stimulus, RED (S1), is also semantically associated with STRAWBERRY (R2), the item studied and tested under the other stimulus (S2). Because Osgood's surface does not consider such potential S1-R2 relationships, it does not address the conditions relevant to the present experiments.

Although such compound category-cue hypotheses can explain the impairment found in Experiment 1, they cannot accommodate the second-order impairment found in Experiments 2 and 4 and the lack of impairment in Experiment 3a. In Experiments 2 and 4, participants may have noticed that Nrp-Similar items like SOUPS-MUSHROOM were also VEGETABLES, but this observation could not have led Rp+ items like GREEN-EMERALD to block Nrp-Similar exemplars because Rp+ items were not members of that implicit category. Thus, no blocking or resource diffusion could have occurred even if participants used both VEGETABLES and SOUPS as cues. One might persist by noting that the related condition contained additional Rp-VEGETABLES and that it might have been this added competition alone that impaired Nrp-Similar items. We tested this hypothesis in Experiment 3a by having participants study the same materials as were studied in Experiments 2 and 4, but by practicing filler categories. If the presence of additional Rp-VEGETABLES had been sufficient to impair the recall of Nrp-Similar items in Experiments 2 and 4, we should have found Nrp-Similar impairment in Experiment 3a, even though Rp+ items were not practiced. That there was no impairment at all in Experiment 3a makes the increased-competition hypothesis untenable and underscores the necessary role of retrieval practice in producing both within-category and cross-category impairment. Thus, paradigm-specific hypotheses based on participants' use of multiple category cues cannot explain the entire pattern of impairment.

Influence of Episodic Retrieval Cues

Another approach to explaining cross-category inhibition in associative terms emphasizes the episodic nature of our retrieval task. That is, our task requires that participants recall not just any red things but only those red things that were studied during the experimental episode. To model this type of recall, we must assume that participants had a representation of the experimental context to which items became associated and with which participants discriminated presented from nonpresented items. The postulation of such representations and a process by which items become associated with them is widely accepted in memory research and is deemed by most researchers to be crucial in explaining episodic recall performance. If participants used a representation of the context as an additional cue during the final test, we must consider how the use of such a cue might affect our assessment of noninhibitory theories.

Suppose that retrieval practice strengthens not only the association between the category and the retrieved exemplar, but also the association between the general experimental context and that practiced item (see, e.g., McGovern, 1964; Raaijmakers & Shiffrin, 1981; Roediger, Stollon, & Tulving, 1977; and Rundus, 1973, for examples of this approach). Because all of the items are associated with the general contextual representation, strengthening exemplars through retrieval practice should impair items in other categories, even though they do not share a category cue with the practiced item. That is, noninhibitory mechanisms, such as occlusion, should cause *generalized episodic competition* between practiced and all nonpracticed items, the effects of which might cause cross-category inhibi-

tion. Indeed, such generalized episodic competition has been advanced in analogous empirical contexts to account for otherwise unexplained effects of cross-cue interference. For example, retroactive interference in the A-B, C-D paradigm, in which stimulus-specific competition cannot occur, has been attributed to competition for a common episodic retrieval cue (McGovern, 1964; see also Postman et al., 1968, for a contrasting account in terms of response suppression). Similar explanations may be offered for the between-category and between-associates impairment found in studies of output interference (Roediger & Schmidt, 1980; Smith, 1971, 1973; Smith, D'Agostino, & Reid, 1970).

Although participants undoubtedly represent the episodic context in which items occur in the retrieval-practice paradigm, competition among Rp+ and Nrp-Similar items for the episodic retrieval cue cannot explain the cross-category inhibition observed in the present studies. Unlike previous paradigms to which this explanation has been applied, the present one involved a within-subjects baseline against which we measured cross-category inhibition. As such, any episodically based competition between Rp+ and Nrp-Similar items should have reduced recall of Nrp-Similar items in both the related and unrelated conditions, eliminating differences in recall. The fact that Nrp-Similar items in the related condition were less well recalled thus implies an additional mechanism by which these items were specifically impaired (see Postman & Underwood, 1973, p. 25, for a discussion of subset-specific response-set suppression in retroactive interference). This finding suggests that other instances of cross-cue impairment that have been attributed to episodic competition might be caused by the processes at work here, a speculation supported by the observation that all such cases have been found in paradigms the effects of which strongly depend on the active retrieval of competitors.

Summary

The preceding discussion illustrates that even relatively complex (but plausible) applications of associative interference theories to our procedure cannot accommodate the present evidence for cross-category inhibition. In addition, numerous other sources of retrieval failure were considered in the discussions of Experiments 1, 2, 3a, and 3b and shown to be inadequate, suggesting that cross-category inhibition reflects a genuine impairment of affected items. Thus, it appears that classical associative interference processes such as occlusion (e.g., McGeoch, 1936; Mensink & Raaijmakers, 1988; Rundus, 1973; Woodworth, 1938), resource diffusion (e.g., J. R. Anderson, 1983), and associative unlearning (e.g., Melton & Irwin, 1940) cannot provide a sufficient account of episodic forgetting and that some additional process must be postulated. In the next section, we discuss how inhibition might be that additional mechanism.

Inhibitory Alternatives

Are there particular inhibitory theories that can accommodate both the cross-category and second-order inhibition effects better than noninhibitory theories? In this section, we outline two inhibitory approaches to retrieval-induced forgetting: lat-

eral inhibition and pattern suppression. Although speculative, these approaches account for both effects and provide a basis for further theoretical development.

Lateral Inhibition

By far the most common and concrete application of inhibition in cognitive models is lateral inhibition. The mechanisms of lateral inhibition, as they are used in computational models of cognitive processes, are based on an analogy to the widespread mechanism of lateral inhibition in the nervous system. Lateral inhibition refers to phenomena whereby one neuron inhibits another neuron, usually through a third (inhibitory) interneuron connecting the two. This form of inhibition is found in both sensory and motor systems, and it can be found at many levels of hierarchical organization within these systems (e.g., see Gallistel, 1980, for a discussion of lateral inhibition in the production of action). That such processes might play a role in the resolution of competition and selection on a cognitive level, particularly in the cases of perception and selective attention, has been argued by several authors (Estes, 1972; Konorski, 1967; McClelland & Rumelhart, 1981; Walley & Weiden, 1973) and appears to be assumed by most advocates of connectionist modeling of cognitive processes.

The straightforward extension of lateral inhibition to the process of memory retrieval is illustrated in Figure 6 in terms of the materials used in Experiment 1. As in the noninhibitory models discussed in the introduction, activation is presumed to spread from the category retrieval cue to the associated exemplars in proportion to the strength of the association linking the category to the exemplar. The new additions in this model are the lateral inhibitory connections (denoted by lines with darkened circles at the end) linking exemplars of a category. Note that such links do not connect members of distinct categories, nor do they extend "vertically" from exemplars to category nodes, representing the assumptions that inhibitory processes (a) are restricted to items bearing an a priori similarity to one another and (b) occur only between members within a single level in the hierarchical system. Essentially, this model has been proposed or at least alluded to by those who have considered the role that lateral inhibition might play in memory retrieval (e.g.,

see Blaxton & Neely, 1983; Martindale, 1981; Roediger & Neely, 1982).

The lateral inhibition model can account for both the within-category and cross-category impairment observed in the retrieval-practice paradigm. During retrieval practice of RED-BLOOD (with the cue RED BL____), activation spreads to all things associated with RED, which then mutually inhibit one another in proportion to their respective levels of activation. Because of the additional letter cues favoring BLOOD, and perhaps because of strengthening accrued by RED-BLOOD on previous practice trials, BLOOD will win the competition, inhibiting TOMATO, which we must assume decreases the likelihood of later recalling TOMATO. Furthermore, because STRAWBERRY is in fact red (in semantic memory), even though it was studied as a FOOD, retrieval practice of RED-BLOOD might affect the later recall of FOOD-STRAWBERRY by a lateral inhibitory connection linking BLOOD to STRAWBERRY (see Figure 6). Because this lateral inhibition affects the representation of STRAWBERRY itself and not just an association with that item, impairment of STRAWBERRY should be evident even though its recall was tested with FOOD. Thus, lateral inhibition predicts cross-category impairment.

The lateral inhibition model also provides an account of the second-order inhibition observed in Experiments 1, 2, and 4. To see this, observe in Figure 6 that lateral inhibitory links connect not only BLOOD, TOMATO, and STRAWBERRY, because of their common membership in the category RED, but also TOMATO, STRAWBERRY, and CRACKERS, because of their common membership in FOOD. When activation spreads from RED to the exemplars TOMATO and STRAWBERRY, those items attempt to inhibit not only BLOOD but also CRACKERS. Because CRACKERS receives no activation from RED, it cannot compete against the inhibition exerted by TOMATO and STRAWBERRY and thus is inhibited even though it bears no direct similarity to BLOOD. This explanation may account for why the replacement of items like GREEN-LETTUCE with nonvegetables such as GREEN-FROG in Experiment 3b eliminated the ability of the retrieval practice of GREEN-EMERALD to inhibit SOUPS-MUSHROOM: Provision of the cue GREEN during retrieval practice no longer activated LETTUCE—a member of VEGETABLES—eliminating the lateral inhibition that would have impaired MUSHROOM. Thus, it is possible for the lateral inhibition model to account for second-order impairment and to explain the dependency of this impairment on similarity to Rp- exemplars.

Two aspects of the present data, however, seem inconsistent with the lateral inhibitory account. The first is the tendency for second-order inhibition to be nearly as great as first-order, within-category impairment. According to the lateral inhibition account, first-order, within-category impairment should be far greater than second-order impairment because retrieval practice of Rp+ items should greatly impair Rp- items, reducing their ability to inhibit Nrp-Similar exemplars. Although the overall trend for the relative impairment of these items was in the correct direction (across Experiments 1, 2, and 4, first-order inhibition averaged 14.9%, whereas second-order inhibition averaged 10.8%), the difference was small, and there were cases in which the second-order inhibition was actually greater than the within-category inhibition. The second finding that seems inconsistent with lateral inhibition is the failure to observe Nrp-




Figure 6. Lateral inhibitory model of items from Experiment 1. Curved solid lines with dark circles at the end indicate lateral inhibitory connections. Straight solid lines indicate studied category-exemplar pairs; heavy lines indicate the subset of those pairs that received retrieval practice; thin dashed lines indicate a preexisting category-exemplar relation that was not studied in any of the experiments.

Dissimilar facilitation in Experiments 2 and 4, even though such facilitation should result from the inhibition of Nrp-Similar items. For example, practice of GREEN-EMERALD impaired SOUPS-MUSHROOM; this inhibition should have enhanced recall of SOUPS-CHICKEN (relative to the recall of SOUPS-CHICKEN in the unrelated condition) by virtue of MUSHROOM's decreased ability to exert lateral inhibition on CHICKEN at the time of the final test. No such facilitation was in evidence. Despite these discrepancies, the lateral inhibition model accounts for the present data more effectively than any of the noninhibitory models discussed previously.

Pattern Suppression

Although less developed computationally than lateral inhibition, pattern suppression may be better able to account for the overall pattern of results. Thus far, all of the mechanisms advanced have assumed processes that operate on discrete, unitary conceptual representations. That is, categories and exemplars are treated as simple units with coherent properties and without internal structure. Another approach would be to use more complex distributed representations of memory items; in fact, we believe that a distributed solution to the problem of retrieval-induced forgetting yields the most natural and intuitive account of our findings. As will be seen, the strength of this approach lies both in the representation of memory items as patterns of semantic features and in the action of a process of pattern suppression.

The basic elements of a distributed approach to retrieval-induced forgetting are illustrated in Figure 7. In this figure, materials from Experiment 2 are represented by sets of semantic feature units, with each unit depicted by a small circle. These units accumulate activation either when the feature is present in the environment or when it is activated by associated units. Larger circles delimit the set of semantic features that represents a given memory item, with features inside a circle tending to mutually activate one another. The intersection of any two large circles contains features shared by those memory items, and activation of these units activates both patterns. Note that the exemplars EMERALD and LETTUCE intersect because of their common membership in the category GREEN; MUSHROOM and

CHICKEN intersect because of their membership in SOUPS; and LETTUCE and MUSHROOM intersect because of their membership in the implicit category VEGETABLES. Note, however, that EMERALD and MUSHROOM do not intersect, reflecting their status as second-order similarities. In this feature-based representation, the retrieval of an item corresponds to the reactivation of all and only those feature units used to represent the item during the study phase (cf. McClelland & Rumelhart, 1986).

The preceding conception of retrieval, together with the tendency for similar items to overlap in their feature units, forms the basis of the way in which this model explains cross-category and second-order inhibition. Consider, for example, what would happen during retrieval practice of GREEN-EMERALD. Presentation of the general category cue GREEN would activate feature units shared by both EMERALD and LETTUCE, in turn activating the remaining features for both of those items. Because selective retrieval of EMERALD requires reactivation of only its own feature units, some mechanism must suppress active features from competing patterns. Thus, retrieval practice should facilitate features associated with the target item (depicted by the darkened feature units of EMERALD), whereas it should impair features associated with competing nontargets (depicted by the white feature units of LETTUCE). Because most of LETTUCE's feature units will have been inhibited, it should be more difficult to reinstate that item's pattern later (i.e., recall that item). Similarly, retrieval of MUSHROOM should also be impaired because it will have had the misfortune of sharing many of LETTUCE's inhibited feature units. Thus, the second-order inhibition of MUSHROOM arises as a consequence of the suppression of the distinctive features of LETTUCE.

The pattern suppression approach provides an intuitive account of the otherwise improbable finding of second-order inhibition and the strong dependency of this finding on the presence of similar Rp- items (Experiment 3b). Both of these effects fall out naturally from the representation of item similarity in terms of common structural components (e.g., feature units) and from the assumption that inhibition resolves pattern competition. This distributed account differs in important ways from lateral inhibition. First, whereas lateral inhibition predicts facilitation of third-order similarities (e.g., SOUPS-CHICKEN), the distributed approach does not: Third-order similarities do not share features with Rp+ or Rp- items, so they should show neither facilitation nor impairment from retrieval practice, a pattern present in our data. Second, whereas lateral inhibition predicts that first-order similarities should be more inhibited than second-order similarities, the present approach predicts that such effects should vary with similarity relations (i.e., with the extent of featural overlap). Suppose, for example, that Rp+ and Rp- items were similar in that 35% of the features of Rp- items overlapped with those of Rp+ items. Retrieval practice should inhibit the distinctive 65% of the features of Rp- items but leave the overlapping 35% highly active. If Nrp-Similar items were unfortunate enough to have 95% of their features overlap with those distinctive features of Rp- patterns, then Nrp-Similar items should actually be more impaired. This follows because 95% of the features of Nrp-Similar items will have been inhibited, as compared with 65% of those in Rp- patterns. This aspect of the model may provide a better account of the relative amount of impairment presently observed for these two




Figure 7. Pattern suppression model illustrated with items from Experiments 2-4. Large circles represent the entire set of features of the item named above the circle; small circles represent individual features. Small dark circles belong to the strengthened item EMERALD; small white circles are features that were suppressed when EMERALD was retrieved.

item types. Although the present account is speculative and many aspects of it remain unspecified, its core assumptions may provide an interesting basis for theoretical and empirical development.

Cue-Independent Forgetting: An Empirical Criterion for Inhibition

The previous discussion of the lateral inhibition and pattern suppression models illustrates how these models localize impairment to changes in the representations of the affected items themselves. This characteristic suggests an empirical dimension that may be generally useful in distinguishing inhibitory and noninhibitory theories and that may be at work in cross-category inhibition: cue-independent forgetting. That is, if a representation is truly inhibited, performance impairments arising from that inhibition should generalize to any cue used to test that item. This property contrasts with the strong cue dependency predicted by theories that attribute impairment to changes in cue-target associations. If cue-target associations are responsible for an item's impairment, impairment should disappear when a distinct retrieval cue is used. For example, impairment of BANANA caused by retrieval practice of FRUIT-ORANGE should disappear when BANANA is tested with the cue MONKEY. Thus, noninhibitory theories are fundamentally theories of cue-dependent forgetting in the sense advocated by Tulving (1974), whereas inhibitory theories are theories of cue-independent forgetting. This observation allows cue independence to serve as a general empirical criterion by which to establish inhibition in memory.

In this article, we have provided evidence that retrieval-induced forgetting is cue independent. However, two issues remain. The first concerns whether the present impairment truly reflects cue-independent forgetting in the sense that any cue is rendered less effective. That cross-category inhibition occurred at all and that this impairment generalized to recall from the implicit category (Experiment 2) argues that inhibited items are rendered less accessible generally. However, more systematic exploration of which cues do and do not reinstate the ability to recall the impaired items is clearly desirable. Second is the issue of whether other varieties of retrieval will also show the cue independence that is characteristic of episodic retrieval-induced forgetting. The strong similarities in the conditions under which retrieval-induced forgetting has been observed in both episodic and semantic retrieval paradigms certainly suggest a common mechanism mediating the effects. Fortunately, the logic underlying the independent probe method reported herein is sufficiently general that the procedure may be adapted to answer questions of cue independence in these and other domains concerned with inhibition.³ In the next section, we outline some of the conditions that appear necessary for retrieval-induced forgetting to occur.

Closing Remarks on the Conditions of Inhibition in Memory Retrieval

Although the present findings clearly favor the conclusion that retrieval involves inhibitory processes, much remains to be specified about the precise conditions necessary for this inhibi-

tion to occur. The existing data suggest that three conditions are necessary to observe inhibition: target retrieval, competitor interference, and recall testing.

Target Retrieval

Several studies of retrieval-induced forgetting have shown that participants must use underspecified cues to retrieve the target item for processing of that target to significantly impair competitors. That is, participants cannot merely identify or study target items but must attempt to retrieve the target from cues that might also match other items. This point is best illustrated by the Blaxton and Neely (1983) study of semantic memory retrieval described in the introduction. In that study, participants were slower to generate a category exemplar from semantic memory (e.g., FRUIT A) if they had just generated four other exemplars from that category on the basis of underspecified retrieval cues (e.g., FRUIT O). Thus, retrieving prior exemplars from semantic memory with incomplete cues caused retrieval-induced forgetting of these final target items. In contrast, participants were faster, not slower, to generate the target exemplars when those prior exemplars were presented intact for speeded naming (e.g., FRUIT ORANGE). Although naming the intact prime items surely requires retrieval of some variety, it is unlikely that such retrieval processes entail the discrimination of those items from the target exemplar. A similar absence of impairment has been demonstrated with a variation of the retrieval-practice paradigm (Anderson et al., 1993) in which participants were provided with intact exemplars during the retrieval-practice phase and asked to recall the category instead (e.g., FR___ORANGE). Thus, if processing of target exemplars is to impair related memory items, targets must be retrieved, and their retrieval cues must be underspecified in a way that leads participants to experience discrimination or response selection difficulties.⁴

³ Indeed, the cue-independence criterion can serve as an empirical criterion for inhibition in nearly any domain. If a representation of any sort is truly inhibited, then attempts to use that representation in behavior should be hindered, regardless of the process or association that makes reference to the representation. That such a strategy for identifying inhibition is indeed general is evidenced by the existence of isomorphs to the independent probe method in research on Pavlovian conditioning. In illustrating that conditioned inhibitors actually inhibit a general representation of an unconditioned stimulus rather than particular stimulus-response associations, conditioned inhibitors are tested in a transfer paradigm in which they are paired with a novel conditioned stimulus (cue) that has been trained with the target unconditioned stimulus thought to be inhibited (Pavlov & Konorski, 1948, as cited in Rescorla, 1975).

⁴ It is crucial to note that apparent exceptions to this rule will arise under well-defined circumstances. For example, if participants study items from a semantic category and then have several members of that category strengthened by their repeated presentation (e.g., FRUIT ORANGE), the remaining target items will be impaired if participants are free to recall exemplars in any order on a final test. The unpracticed targets will be impaired not because of the repeated presentation of practiced items but because the heightened availability of those practiced items will lead them to be reported early on the final test, causing retrieval-induced forgetting on unpracticed targets. Failure to control for output interference of this nature may be the single most important reason why occlusion models have retained their popularity (see Anderson, Bjork, & Bjork, 1994, for this argument).

Competitor Interference

Although target retrieval appears necessary for inhibition to occur, it is not by itself sufficient to cause impairment. Whether the generation of a target item impairs related representations depends on whether those related items interfere with the retrieval of the target. This dependency has been illustrated in several ways with the retrieval-practice paradigm. Consider, for example, the results of an experiment conducted by Anderson et al. (1994). In that experiment, participants studied categories that were composed of all strong exemplars (i.e., high taxonomic frequency items like FRUIT ORANGE), all weak exemplars (i.e., low taxonomic frequency items like FRUIT GUAVA), or half strong and half weak exemplars. Anderson et al. reasoned that if inhibition reduces interference from strong competitors, then retrieval practice should have very different effects on competitors as a function of whether they are strong or weak exemplars. The expected dependency was observed: Whereas retrieval practice always impaired the later recall of strong exemplars, such practice left the later recall of weak exemplars unaffected, despite significant facilitation of practiced targets. These findings demonstrate that the retrieval of memory targets from underspecified cues is not enough to impair similar items; rather, impairment hinges on whether similar competitors actually interfere with the retrieval of target items (see the *Results and Discussion* section of Experiment 3b for a related discussion in the context of those findings; see also Carr & Dagenbach, 1990, and Simpson & Kang, 1994, for related findings in semantic memory).⁵ The lack of impairment of weak exemplars also demonstrates how strengthening of practiced items resulting from retrieval practice does not impair all related items, contradicting what one might expect on the basis of models that attribute impaired recall to the intrusion of dominant competitors (McGeoch, 1936; Mensink & Raaijmakers, 1988; Rundus, 1973; Woodworth, 1938).

Recall Testing

Studies using procedures related to the retrieval-practice paradigm, such as retroactive interference and part-set cuing, have demonstrated that manipulations that impair recall of items often do not impair recognition accuracy for the same items (e.g., Postman & Stark, 1969; Slamecka, 1975; Todres & Watkins, 1981). For example, in a classic study conducted by Postman and Stark (1969), significant retroactive interference in an A-B, A-C paradigm was completely eliminated when a multiple-choice recognition test rather than the traditional cued recall procedure was used to assess participants' memory of A-B responses. Although retroactive interference could differ from the present phenomenon, use of the method of anticipation in interference paradigms may add a significant retrieval-induced forgetting component to many cases of retroactive interference. Given this argument, the observation that retroactive interference dissipates on presentation of intact study items during a recognition accuracy test implies that similar results may be obtained with paradigms specifically directed at the study of retrieval-induced forgetting. Thus, the consequences of inhibitory mechanisms might not be observable if inhibited responses are provided intact to the participant (with the possible excep-

tion of recognition latency; see Neely, Schmidt, & Roediger, 1983). This dissipation of impairment is consistent with the notion that the representations of items are inhibited by retrieval but not damaged in any permanent sense (see Postman & Stark, 1969, for a related point).

Selective Retrieval as a Problem of Selective Attention

The discussion of inhibitory mechanisms has thus far concentrated on domain-specific issues regarding the status of inhibition in theories of memory retrieval. We now return to a discussion of how our findings may be of relevance to those not directly concerned with retrieval processes. Does the present evidence for inhibition reflect mechanisms peculiar to retrieval, or might our findings result from a more general inhibitory mechanism common to a variety of cognitive functions?

Although we do not exclude the former, more restricted interpretation of our findings, we offer the broader view that the mechanisms at work during the sort of retrieval in which our participants were engaged—that is, selective retrieval of particular target items—are analogous to, or perhaps the same as, those mechanisms at work when participants must selectively attend to objects in the external environment. In this section, we first set out the computational and empirical similarities between selective retrieval and selective attention. We then discuss issues in extending attentional inhibitory mechanisms to the domain of episodic memory retrieval. We end by briefly noting the correspondence of the proposed interpretation of our findings to neuropsychological hypotheses about the substrates of selective retrieval.

Functional and Empirical Similarities Between Retrieval and Attention

Few people would contest the proposition that attention can be shifted from objects in the external world to "objects" in the internal world—images, facts, and episodes generated on the basis of past experience. For example, to determine what one had for dinner last evening requires that the reader cease, for the moment, reading this article and refocus attention to the internal representation of last evening. Nonetheless, contemporary research on attention has centered almost exclusively on attention directed to external percepts. To highlight this somewhat neglected domain of attention, we introduce the term *con-*

⁵ The notion that inhibitory processes are invoked to overcome interference from memory competitors is consistent with the assumption, common in classical interference theory, that the degree of retroactive interference attributable to unlearning is related to the amount of negative transfer observed on the acquisition of interpolated study lists. For instance, Melton and Irwin (1940) viewed the occurrence of both overt and covert intrusions during the acquisition of interpolated lists as a necessary condition for the unlearning of those intruding responses. This view, which has been called the elicitation hypothesis (e.g., see Postman & Underwood, 1973), has received some direct empirical support; manipulations of variables that either increase or decrease the expected frequency of covert intrusions during the learning of interpolated lists affect the amount of retroactive interference observed in the expected manner (Birnbau, 1968; Friedman & Reynolds, 1967; Goggin, 1967; Postman, Keppel, & Stark, 1965).

ceptually focused selective attention to refer to instances in which focal awareness is directed to some specific internal representation, the referent of which is no longer present in the environment. In this section, we consider the merits of viewing selective retrieval as an instance of conceptually focused selective attention, arguing that such a view links the present evidence for inhibition with analogous inhibitory findings in the literature on *perceptually focused selective attention* and crystallizes the emerging trend to expand the domain of attention to include both conceptual and perceptual foci (Neill et al., in press; Neill & Westberry, 1987; Neumann & DeSchepper, 1992; Posner, 1980, 1987).

Although memory retrieval and selective attention are typically seen as separate research topics, similarities in the computational problems to be solved in these cognitive functions suggest that it may be worth reconsidering the division. Consider, for example, Figure 8, which illustrates schematically the problems of perceptually focused selective attention and selective retrieval.⁶

On the left side of Figure 8, the representations of several fruits receive direct activation from perceptual input, and the function of selective attention ("A" in the figure) is to isolate one representation among the many competitors activated by the senses. On the right side of Figure 8, the representations of several fruits are activated by an externally or internally generated retrieval cue, and the function of retrieval is to isolate one representation among the many competitors activated by that cue. The principal differences between selective retrieval and selective attention are that, in selective retrieval, competition among active representations is conceptually initiated rather than perceptually initiated and the output of the retrieval mechanism is a consciously experienced memory rather than a consciously experienced percept. In both cases, proper task performance usually necessitates selection of a single representation

from among the activated alternatives and rejection of the remainder. If one regards this function of selection as basic to the task of perceptually focused selective attention, then selective retrieval should be viewed as conceptually focused selective attention (for related perspectives, see Dagenbach et al., 1990; Gernsbacher, 1991; Gernsbacher & Faust, 1991; Neill et al., in press; Neill & Westberry, 1987; Neumann & DeSchepper, 1991).⁷

If one views selective retrieval as a problem of selective attention, the present findings suggest that conceptually focused and perceptually focused attention involve similar mechanisms of selection. In particular, the present findings suggest that conceptual attention operates both by facilitating the to-be-attended memory targets and by actively suppressing the to-be-ignored, interfering competitors. This dual process view parallels recent analyses of perceptual attention in which inhibition is thought to underlie the rejection of to-be-ignored percepts. For instance, a flood of recent work on the phenomenon of *negative priming* (Tipper, 1985) has argued that the act of ignoring a percept may be achieved through active inhibition. In an early study conducted by Tipper (1985), participants viewed a line drawing of a common object (e.g., a hammer), colored in red, superimposed over a second drawing, colored in green. Participants were instructed to name the stimulus in red and to ignore the stimulus in green. When the stimulus in red had just been ignored on the previous trial (i.e., the stimulus had appeared in green), participants were slower to name that red stimulus than when an unrelated picture had been previously ignored. Tipper interpreted the longer response latencies to name previously ignored objects as evidence that participants had actively inhibited the interfering, to-be-ignored stimulus, the consequences of which became evident only when the ignored stimulus suddenly became task relevant on the next trial. The negative priming effect, which has since been observed across a broad range of stimulus and task conditions (see Neill et al., in press, for a re-

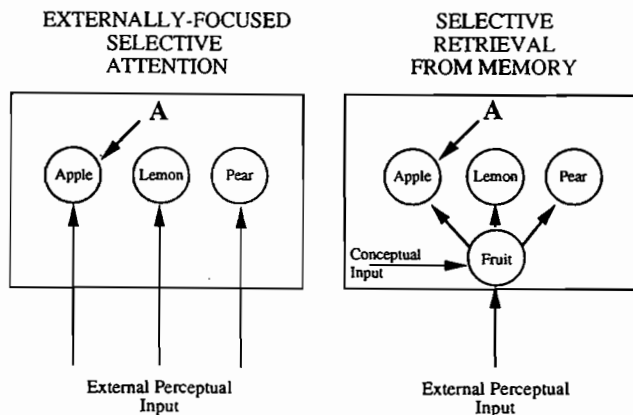


Figure 8. Schematic diagram of the relationship between externally focused selective attention and selective memory retrieval. The circumstance of externally focused selective attention is depicted on the left, and that of selective memory retrieval is depicted on the right. Circles designate concepts for the terms named; arrows from outside the box designate perceptual input; arrows from inside the box represent internal sources of activation either from concepts serving as retrieval cues or from attention (A).

⁶ The present discussion assumes a late-selection view of selective attention (Deutsch & Deutsch, 1963). In the late-selection view of attention, all objects in the environment are first semantically identified and only then are selected to be (or not to be) the objects of focal awareness. Although this view has many proponents (Allport, Tipper, & Chmiel, 1985; Duncan, 1980; Keele & Neill, 1978; Marcel, 1983; Neill, 1989; Norman, 1968; Posner, 1982; Schneider & Shiffrin, 1977), it is not universally accepted (see Allport, 1989, for a discussion). However, because the empirical findings to which we compare the present memory results favor a late-selection view of attention, it is the one we adopt in building the present analogy.

⁷ We do not intend the present argument to imply that all cases in which the nervous system resolves selection problems should be regarded as problems of attention. For instance, most authors do not believe that the thousands of force, direction, and speed parameter selections that occur in the execution of a simple motor skill are resolved by conscious attentional mechanisms. Concern over what distinguishes instances of selection requiring attention from those that do not is warranted, but this concern is not unique to the present argument. We maintain, however, that selection is an important function of attention (uncontroversially, we assume); because the present retrieval-practice task clearly requires both selection and attention (as our participants would most assuredly attest) and because the notion of conceptual selective attention seems subjectively sound, we believe that the present analogy is useful.

view), suggests that conceptually and perceptually focused selective attention may share similar inhibitory principles of selection.

Although retrieval-induced forgetting and negative priming share functional and empirical similarities, it remains unclear whether these findings are better viewed as stemming from identical or isomorphic attention processes (cf. Posner, Inhoff, Friedrich, & Cohen, 1987). For instance, the highly perceptual nature of Tipper's (1985) object recognition procedure immediately suggests special visual attention mechanisms acting on representations of objects. Such mechanisms may be functionally and anatomically distinct from those used in selective retrieval paradigms. Several findings suggest, however, that negative priming may not be tied to perceptual representations. Tipper (1985) found that ignoring a line drawing (e.g., of a cat) also slowed the subsequent naming of a semantically related line drawing (e.g., of a dog), although to a lesser extent. Similar results have been observed with verbal associates (Neumann et al., 1993; Yee, 1991) and with mixtures of pictures and words (Tipper & Driver, 1988), suggesting that ignored percepts are processed to the level of semantics and then inhibited. Accordingly, the increased latency to name visual objects in negative priming could be characterized as a special case of retrieval inhibition in semantic memory (Blaxton & Neely, 1983; Brown, 1979, 1981; Brown et al., 1985) in which interfering competitors become active directly by external cues rather than indirectly by retrieval cues. To the extent that negative priming affects semantic representations of ignored objects, a common inhibitory mechanism mediating perceptual and conceptual selective attention attains plausibility.

Even if negative priming and retrieval-induced forgetting derive from distinct attentional systems, however, the analogy between perceptual and conceptual selective attention may still be of value. For instance, Posner (1980, 1987) argued that an understanding of visual selective attention might serve as a useful model by which to understand internally focused attention. Although few explicit attempts have been made to relate perceptual and conceptual attention in this manner, recent extensions of negative priming to domains outside of visual attention may be viewed in these terms. For instance, Neumann and DeSchepper (1992) showed negative priming effects in a short-term memory paradigm, demonstrating that selective attention mechanisms may be applied to memorial representations long after perceptual input has ceased. Similarly, Simpson and Kang (1994) demonstrated that negative priming effects occur when people resolve the meanings of ambiguous words (e.g., *bank*, with senses *river* and *money*) in favor of a context-appropriate meaning. Extending negative priming in this manner to non-perceptual tasks may exemplify exactly the sort of transfer between perceptual and conceptual attention research that Posner anticipated. The present conception of selective retrieval as conceptual selective attention makes this transfer explicit, crystallizing what must already be assumed to interpret the preceding phenomena as negative priming. In the next section, we discuss specific issues arising in the extension of this general functional analysis to our episodic retrieval findings.

Episodic Forgetting as Evidence for Attentional Inhibition

If our functional analysis of the relationship between selective retrieval and selective attention is correct, then finding retrieval-

induced forgetting in episodic memory should not seem surprising. Nonetheless, the durability of episodic forgetting might lead some to conclude that the effect cannot result from inhibitory mechanisms and, thus, that the mechanisms of episodic memory impairment must differ from those at work in attentional paradigms. Our findings would then have significant implications for associative theories of episodic forgetting but would not have broader implications for the role of inhibition in other cognitive processes. In this section, we address these concerns, arguing two main points: (a) that the durability of episodic forgetting should not be taken as evidence that inhibitory processes are not at work and (b) that the particular differences in the time course of impairment in episodic memory and selective attention paradigms may reflect variations in the impact of inhibition on different representations rather than differences in the inhibition process itself. Thus, episodic forgetting may have characteristics that are limited to episodic representations and still reflect the function of general mechanisms of attentional inhibition.

Discrepancy With Physiological Inhibition

The results of Experiment 4, as well as analogous findings observed by Anderson et al. (1994), show that retrieval-induced forgetting endures for at least 20 min. One might be reluctant to attribute these and other such enduring effects to inhibitory processes on the grounds of the rather large discrepancy between the durability of these effects and those typically observed at the level of neurophysiology (the effects of which typically last between a few milliseconds and a few seconds; see, e.g., Kandel & Schwartz, 1991).

Although this perspective may seem reasonable at first, it makes the strong assumption that cognitive inhibition should follow the characteristics of individual neurons receiving a single inhibitory input, an assumption that is likely to be far too simple. For example, for these assumptions to be accurate, one must extrapolate the temporal characteristics of neuronal inhibition to the more complex units of organization that surely underlie cognitive tasks. Because the neurophysiology of such complex circuitry is sufficiently unclear for one to exclude inhibitory effects with different temporal parameters, the timing of neuronal inhibition should not seriously constrain hypothesized characteristics of cognitive inhibition. Indeed, were one to constrain what was and was not possible in more complex learning situations on the basis of neuronal physiology (e.g., long-term potentiation), one would be forced to conclude that human beings were incapable of retaining information for longer than a week, clearly an unwarranted conclusion.

Even if one could assume that the temporal parameters of inhibition in complex systems were the same as those exhibited neuronally, the present phenomenon might still be the result of inhibitory mechanisms. For example, nothing in the inhibitory approach to memory retrieval requires that inhibition be both the cause and the effect of retrieval-induced forgetting. That is, it is possible for retrieval-based inhibitory processes to cause a separate disruptive effect on an episodic representation. Depending on the characteristics of this disruption, impairment of any duration may be possible. To see this, consider how an inhibitory process might disrupt a set of synchronously firing

computational units. Even if the local effects of inhibitory input into a single unit in the set are brief, the disruption of the global synchrony of all of the units might endure. Any later task that relied on the synchronous firing of those units could thus be disrupted without that disruption reflecting a persisting inhibition in a neuronal sense. Many such accounts of enduring impairment are possible.⁸ Given these possibilities, the discrepancy between the durability of the present effects and those observed physiologically should not lead one to exclude physiological inhibitory mechanisms as the underlying cause of the effect.

Discrepancy With Attentional Inhibition

If the inhibitory mechanisms used during episodic retrieval can have such an enduring impact, why are the effects of attentional inhibition not similarly enduring? For instance, negative priming effects typically last for only several hundred milliseconds after the stimulus item has been ignored (although exceptions to this have been found; see, e.g., Tipper, Weaver, Cameron, Brehaut, & Bastedo, 1991). Although one might see this difference as evidence that the processes responsible for impairment in the two cases must be fundamentally different, a common inhibitory process may in fact underlie effects of varying duration if the representations involved are differentially affected by inhibition.

Consider how the representations underlying memory retrieval and selective attention tasks differ. Selective attention procedures typically demand access to a general semantic representation of the target stimulus thought to be inhibited. For example, the negative priming task described earlier requires participants to name pictures, a task that we assume demands access to a semantic representation of that item. Semantic representations, particularly of the sort that underlie the simple objects and letters used in attention experiments, are highly overlearned and may be resistant to any long-term effects of inhibitory processes. In contrast, episodic memory paradigms, such as the present one, require access to a unique episodic representation of an item formed on a single occasion in a brief study phase. Such representations are likely to be more vulnerable to any long-term disruptive effects of inhibitory processes. Indeed, insofar as the disruptive effects of inhibition occur to the unique spatiotemporal aspects of episodic encodings, semantic representations may suffer no disruption at all, their impairment merely reflecting the fleeting effects of the initial inhibitory inputs to the interfering items.⁹

The important observation is that the discrepancy in the durability of the impairment in episodic memory and attentional paradigms may in fact stem from a difference in the properties of the representations involved in the two cases. Although the characteristics of the impairment resulting from inhibition may be quite different for episodic and semantic representations, it is entirely possible that a common inhibitory process underlies these effects. Granting this possibility, the present data suggest the intriguing notion that people's experience of long-lasting episodic forgetting may be directly linked to the very mechanisms that allow them to selectively attend both to the external world and to internal representations.

From Behavior to Brain

Although the present investigation of inhibitory mechanisms has been purely behavioral, our conclusions regarding the inhibitory basis of selective retrieval correspond nicely to implications derivable from neuropsychological accounts of strategic memory retrieval. Several authors have noted that directed memory retrieval tasks have a strong strategic component that may be controlled by the prefrontal cortex (Baddely & Wilson, 1986; Goldberg & Bilder, 1986; Goldberg & Costa, 1986; Luria, 1976; Mercer, Wapner, Gardner, & Benson, 1977; Moscovitch, 1989; Shapiro, Alexander, Gardner, & Mercer, 1981; Shimamura, in press; Stuss & Benson, 1986; Tulving, 1989). In particular, on the basis of observations of frontal patients' extreme vulnerability to memory interference and memory perseveration, the prefrontal cortex is thought to be important in (a) the resolution of competition arising when the memory system delivers too many (or contextually inappropriate) memory traces in response to a retrieval cue (Moscovitch, 1989) and (b) the prevention of highly activated memory competitors from repeatedly capturing control of retrieval mechanisms (Luria, 1976). Thus, the strategic memory functions attributed to the prefrontal cortex correspond to the functions thought to be served by the presently hypothesized attentional inhibitory mechanisms.

Although the mechanism by which the prefrontal cortex resolves memory competition remains unclear, it is worth noting that this brain region has long been associated with general executive-control functions, including the use of inhibition to reduce perceptually initiated task interference (Diamond, 1990; Fuster, 1989; Luria, 1966; Mishkin, 1964; Shallice, 1988). If executive inhibitory functions resolve perceptually initiated interference on nonmemory tasks, and if the aforementioned authors' speculations about strategic retrieval are correct, it seems reasonable to suppose that executive inhibitory processes may also resolve conceptually initiated interference during selective retrieval, as we have proposed here. Thus, the present empirical demonstration of inhibitory processes in selective retrieval may be seen as a clear instance of the normal functioning of inhibi-

⁸ Another approach to explaining durable inhibition would be to assume sustained or tonic inhibitory input to the impaired items, perhaps stemming from the sustained facilitation on practiced exemplars. If inhibitory input could be continuously applied to interfering items, inhibitory effects of nearly any duration may be possible. We do not favor this account, however, because it presumes a strong link between strengthening of practiced items and impairment.

⁹ Another factor accounting for differences in characteristics of impairment observed in selective attention and episodic retrieval paradigms may be the sensitivity of the tests used to measure inhibition in the two domains; whereas negative priming studies have universally presented the intact stimulus to participants for speeded identification, the retrieval-practice paradigm uses the more difficult task of cued retrieval. As noted in our previous discussion of the conditions of inhibition, memory tests (e.g., recognition) that provide participants the intact stimulus that is thought to be inhibited virtually eliminate retroactive interference and part-set cuing inhibition in long-term memory (Postman & Stark, 1969; Slamecka, 1975; Todres & Watkins, 1981). Consequently, the time course of attentional inhibition may be underestimated by existing paradigms.

tory processes long suspected to be deficient in frontal patients; the extreme inability of people suffering from such damage to maintain the focus of behavior and thought may thus provide a vivid reminder of the importance of inhibitory control in cognition.

CONCLUDING REMARKS

Proponents of the use of inhibitory constructs in theories of cognition often experience lingering uncertainty regarding the empirical status of cognitive inhibition. One reason for this uncertainty stems from the fact that most performance deficits that might plausibly be attributed to inhibitory processes can also be explained with theories that use only excitatory mechanisms. Because of this apparent indeterminacy, discussion of inhibitory alternatives is often abbreviated in the interest of theoretical parsimony. The present article has addressed this empirical indeterminacy of inhibition in the context of retrieval from long-term memory.

The present experiments lead to the conclusion that retrieval of information from long-term memory not only involves the spread of activation from retrieval cues to the desired target but may often require the active inhibition of competitors that interfere with selective retrieval. The consequence of this inhibition is a persisting deficit in recall performance for inhibited items; this deficit in recall appears to generalize to the testing of those items through other retrieval cues and also to items that are merely similar to the suppressed information. The observation of such generalized impairment implies the existence of a form of cue-independent forgetting that has been shown to be inconsistent with modern theories of memory retrieval but that is consistent with what may be part of a general computational solution to the problem of selective attention: inhibition.

References

- Allport, A. (1989). Visual attention. In M. I. Posner (Ed.), *Foundations of cognitive science* (pp. 631–682). Cambridge, MA: MIT Press.
- Allport, A., Tipper, S. P., & Chmiel, N. (1985). Perceptual integration and postcategorical filtering. In M. I. Posner & O. S. M. Marin (Eds.), *Attention and performance* (Vol. 11, pp. 107–132). Hillsdale, NJ: Erlbaum.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, M. C., & Bjork, R. A. (1994). Mechanisms of inhibition in long-term memory: A new taxonomy. In D. Dagenbach & T. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 265–325). San Diego, CA: Academic Press.
- Anderson, M. C., Bjork, E. L., & Bjork, R. A. (1993, November). *Strengthening is not enough: Evidence against the blocking theory of retrieval inhibition*. Paper presented at the 34th Annual Meeting of the Psychonomic Society, Washington, DC.
- Anderson, M. C., Bjork, R. A., & Bjork, E. L. (1994). Remembering can cause forgetting: Retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 1063–1087.
- Baddely, A. D., & Wilson, B. (1986). Amnesia, autobiographical memory, and confabulation. In D. Rubin (Ed.), *Autobiographical memory* (pp. 225–252). Cambridge, England: Cambridge University Press.
- Barnes, J. M., & Underwood, B. J. (1959). "Fate" of first-list associations in transfer theory. *Journal of Experimental Psychology*, 58, 95–105.
- Beech, A., Powell, T., McWilliams, J., & Claridge, G. (1989). Evidence of reduced "cognitive inhibition" in schizophrenia. *British Journal of Clinical Psychology*, 28, 110–116.
- Birnbaum, I. M. (1968). Response class similarity and first-list recall with mixed and unmixed transfer designs. *Journal of Experimental Psychology*, 77, 542–546.
- Bjork, R. A. (1989). Retrieval inhibition as an adaptive mechanism in human memory. In H. L. Roediger III & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 309–330). Hillsdale, NJ: Erlbaum.
- Blaxton, T. A., & Neely, J. H. (1983). Inhibition from semantically related primes: Evidence of a category-specific inhibition. *Memory & Cognition*, 11, 500–510.
- Broadbent, D. E. (1958). *Perception and communication*. London: Pergamon Press.
- Brown, A. S. (1979). Priming effects in semantic memory retrieval processes. *Journal of Experimental Psychology: Human Learning and Memory*, 5, 65–77.
- Brown, A. S. (1981). Inhibition in cued retrieval. *Journal of Experimental Psychology: Human Learning and Memory*, 7, 204–215.
- Brown, A. S., Whiteman, S. L., Cattoi, R. J., & Bradley, C. K. (1985). Associative strength level and retrieval inhibition in semantic memory. *American Journal of Psychology*, 98, 433–447.
- Carr, T. H., & Dagenbach, D. (1990). Semantic priming and repetition priming from masked words: Evidence for a center-surround attentional mechanism in perceptual recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 341–350.
- Cohen, J. D., & Servan-Schreiber, D. (1992). Context, cortex and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review*, 99, 45–77.
- Dagenbach, D., & Carr, T. (Eds.). (1994a). *Inhibitory processes in attention, memory, and language*. San Diego, CA: Academic Press.
- Dagenbach, D., & Carr, T. (1994b). Inhibitory processes in perceptual recognition: Evidence for a center-surround attentional mechanism. In D. Dagenbach & T. Carr (Eds.), *Inhibitory processes in attention, memory and language* (pp. 327–358). San Diego, CA: Academic Press.
- Dagenbach, D., Carr, T. H., & Barnhardt, T. M. (1990). Inhibitory semantic priming of lexical decisions due to failure to retrieve weakly activated codes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 328–340.
- Dempster, F. N. (1991). Inhibitory processes: A neglected dimension of intelligence. *Intelligence*, 15, 157–173.
- Dempster, F. N., & Brainerd, C. J. (Eds.). (1994). *New perspectives on interference and inhibition*. New York: Academic Press.
- Deutsch, J. A., & Deutsch, D. (1963). Attention: Some theoretical considerations. *Psychological Review*, 70, 80–90.
- Diamond, A. (1990). Developmental time course in human infants and infant monkeys and the neural bases of inhibitory control in reaching. In A. Diamond (Ed.), *The development and neural bases of higher cognitive functions* (pp. 637–676). New York: New York Academy of Sciences.
- Duncan, J. (1980). The locus of interference in the perception of simultaneous stimuli. *Psychological Review*, 87, 272–300.
- Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, 36, 433–458.
- Easton, P., & Gordon, P. E. (1984). Stabilization of Hebbian neural nets by inhibitory learning. *Biological Cybernetics*, 51, 1–9.
- Estes, W. K. (1972). An associative basis for coding and organization in memory. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory* (pp. 161–191). New York: Wiley.
- Friedman, M. J., & Reynolds, J. H. (1967). Retroactive inhibition as a

- function of response-class similarity. *Journal of Experimental Psychology*, 74, 351-355.
- Fuster, J. M. (1989). *The prefrontal cortex: Anatomy, physiology, and neuropsychology of the frontal lobe*. New York: Raven Press.
- Gallistel, C. R. (1980). *The organization of action: A new synthesis*. Hillsdale, NJ: Erlbaum.
- Gernsbacher, M. A. (1991). Cognitive processes and mechanisms in language comprehension: The structure building framework. In G. H. Bower (Ed.), *The psychology of learning and motivation* (pp. 217-264). New York: Academic Press.
- Gernsbacher, M. A. (1990). *Language comprehension as structure building*. Hillsdale, NJ: Erlbaum.
- Gernsbacher, M. A., & Faust, M. E. (1991). The mechanism of suppression: A component of general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 245-262.
- Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, 91, 1-67.
- Goggin, J. (1967). First-list recall as a function of second-list learning method. *Journal of Verbal Learning and Verbal Behavior*, 6, 423-427.
- Goldberg, E., & Bilder, R. M. (1986). Neuropsychological perspectives: Retrograde amnesia and executive deficits. In L. Poon (Ed.), *Handbook for clinical memory assessment of older adults* (pp. 55-68). Washington, DC: American Psychological Association.
- Goldberg, E., & Costa, L. D. (1986). Qualitative indices in neuropsychological assessment: An extension of Luria's approach to executive deficit following prefrontal lesions. In I. Grant & K. M. Adams (Eds.), *Neuropsychological assessment of neuropsychiatric disorders* (pp. 48-64). New York: Oxford University Press.
- Hasher, L., & Zacks, R. T. (1988). Working memory, comprehension and aging: A review and a new view. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, pp. 193-225). New York: Academic Press.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295-355.
- Kahneman, D. (1973). *Attention and effort*. Englewood Cliffs, NJ: Prentice Hall.
- Kandel, E. P., & Schwartz, J. H. (1991). *Principles of neural science*. New York: Elsevier.
- Keele, S. W., & Neill, W. T. (1978). Mechanisms of attention. In E. C. Carterette & M. P. Friedman (Eds.), *Handbook of perception* (Vol. 9, pp. 3-47). New York: Academic Press.
- Konorski, J. (1967). *Integrative activity of the brain: An interdisciplinary approach*. Chicago: University of Chicago Press.
- Lewandowsky, S. (1991). Gradual unlearning and catastrophic interference: A comparison of distributed architectures. In W. E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays in honor of Bennet B. Murdock* (pp. 445-476). Hillsdale, NJ: Erlbaum.
- Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, 91, 295-327.
- Luria, A. R. (1966). Disturbances of higher cortical functions with lesions of the frontal regions. In *Higher cortical functions in man* (pp. 218-295). London: Tavistock.
- Luria, A. R. (1976). *The neuropsychology of memory*. Washington, DC: Wiley.
- Marcel, A. J. (1983). Conscious and unconscious perception: An approach to the relations between phenomenal experience and perceptual processes. *Cognitive Psychology*, 15, 238-300.
- Martindale, C. (1981). *Cognition and consciousness*. Homewood, IL: Dorsey Press.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, 88, 375-407.
- McClelland, J. L., & Rumelhart, D. E. (1986). A distributed model of human learning and memory. In J. L. McClelland & D. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2: Psychological and biological models* (pp. 170-215). Cambridge, MA: MIT Press.
- McDowd, J. M., Oseas-Kreger, D. M., & Fillion, D. L. (1994). Inhibitory processes in selective attention and aging. In F. N. Dempster & C. J. Brainerd (Eds.), *New perspectives on interference and inhibition in cognition*. New York: Academic Press.
- McGeoch, J. A. (1936). Studies in retroactive inhibition: VII. Retroactive inhibition as a function of the length and frequency of presentation of the interpolated lists. *Journal of Experimental Psychology*, 19, 674-693.
- McGovern, J. B. (1964). Extinction of associations in four transfer paradigms. *Psychological Monographs*, 78, (16, Whole No. 593).
- Melton, A. W., & Irwin, J. M. (1940). The influence of degree of interpolated learning on retroactive inhibition and the overt transfer of specific responses. *American Journal of Psychology*, 3, 173-203.
- Mensink, G. J. M., & Raaijmakers, J. G. W. (1988). A model of interference and forgetting. *Psychological Review*, 95, 434-455.
- Mercer, B., Wapner, W., Gardner, H., & Benson, D. F. (1977). A study of confabulation. *Archives of Neurology*, 34, 429-433.
- Mishkin, M. (1964). Perseveration of central sets after frontal lesions in monkeys. In J. M. Warren & K. Akert (Eds.), *The frontal granular cortex and behavior* (pp. 219-237). New York: McGraw-Hill.
- Moscovitch, M. (1989). Confabulation and the frontal systems: Strategic versus associated retrieval in neuropsychological theories of memory. In H. L. Roediger & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 133-160). Hillsdale, NJ: Erlbaum.
- Neely, J. H., Schmidt, S. R., & Roediger, H. L. III. (1983). Inhibition from related primes in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9, 196-211.
- Neill, W. T. (1989). Ambiguity and context: An activation-suppression model. In D. S. Gorfein (Ed.), *Resolving semantic ambiguity* (pp. 63-83). New York: Springer-Verlag.
- Neill, W. T. (1991, August). *Consciousness and the inhibitory control of cognition*. Paper presented at the 99th Annual Convention of the American Psychological Association, San Francisco, CA.
- Neill, W. T., Valdes, L. A., & Terry, K. M. (in press). Selective attention and the inhibitory control of cognition. In F. N. Dempster & C. J. Brainerd (Eds.), *New perspectives on interference and inhibition in cognition*. New York: Academic Press.
- Neill, W. T., & Westberry, R. L. (1987). Selective attention and the suppression of cognitive noise. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 327-334.
- Neumann, E., Cherau, J. F., Hood, K. L., & Steinnagel, S. L. (1993). Does inhibition spread in a manner analogous to spreading activation. *Memory*, 1, 81-105.
- Neumann, E., & DeSchepper, B. G. (1991). Costs and benefits of distractor inhibition in selective attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 1136-1145.
- Neumann, E., & DeSchepper, B. G. (1992). An inhibition-based fan effect: Evidence for an active suppression mechanism in selective attention. *Canadian Journal of Psychology*, 46, 11-50.
- Norman, D. A. (1968). Toward a theory of memory and attention. *Psychological Review*, 75, 522-536.
- Norman, D. A., & Bobrow, D. G. (1975). On data-limited and resource-limited processes. *Cognitive Psychology*, 7, 44-64.
- Norman, D. A., & Shallice, T. (1980). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation: Advances in research and theory* (Vol. 4, pp. 1-18). New York: Plenum.
- Osgood, C. E. (1949). The similarity paradox in human learning: A resolution. *Psychological Review*, 56, 132-143.

- Park, J., & Kanwisher, N. (1994). Negative priming for spatial locations: Identity mismatching, not distractor inhibition. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 613-623.
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32A, 3-25.
- Posner, M. I. (1982). Cumulative development of attentional theory. *American Psychologist*, 37, 168-179.
- Posner, M. I. (1987). Selective attention and cognitive control. *Trends in Neuroscience*, 10, 13-17.
- Posner, M. I., Inhoff, A. W., Friedrich, F. J., & Cohen, A. (1987). Isolating attentional systems: A cognitive-anatomical analysis. *Psychobiology*, 15, 107-121.
- Postman, L., Keppel, G., & Stark, K. (1965). Unlearning as a function of the relationship between successive response classes. *Journal of Experimental Psychology*, 69, 111-118.
- Postman, L., & Stark, K. (1969). The role of response availability in transfer and interference. *Journal of Experimental Psychology*, 79, 168-177.
- Postman, L., Stark, K., & Fraser, J. (1968). Temporal changes in interference. *Journal of Verbal Learning and Verbal Behavior*, 7, 672-694.
- Postman, L., & Underwood, B. J. (1973). Critical issues in interference theory. *Memory & Cognition*, 1, 19-40.
- Raaijmakers, J. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, 88, 93-134.
- Ratcliff, R., Clark, S. E., & Shiffrin, R. M. (1990). The list-strength effect: I. Data and discussion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 163-178.
- Rescorla, R. A. (1975). Pavlovian excitatory and inhibitory conditioning. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes* (Vol. 2, pp. 7-36). Hillsdale, NJ: Erlbaum.
- Roediger, H. L. III, & Neely, J. H. (1982). Retrieval blocks in episodic and semantic memory. *Canadian Journal of Psychology*, 36, 213-242.
- Roediger, H. L. III, & Schmidt, S. R. (1980). Output interference in the recall of categorized and paired-associate lists. *Journal of Experimental Psychology: Human Learning and Memory*, 6, 91-105.
- Roediger, H. L. III, Stellan, C. C., & Tulving, E. (1977). Inhibition from part-list cues and rate of recall. *Journal of Experimental Psychology: Human Learning and Memory*, 3, 174-188.
- Rundus, D. (1973). Negative effects of using list items as retrieval cues. *Journal of Verbal Learning and Verbal Behavior*, 12, 43-50.
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search and attention. *Psychological Review*, 84, 1-66.
- Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge, England: Cambridge University Press.
- Shapiro, B. E., Alexander, M. P., Gardner, H., & Mercer, B. (1981). Mechanisms of confabulation. *Neurology*, 31, 1070-1076.
- Shimamura, A. P. (in press). Memory and frontal lobe function. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences*. Cambridge, MA: MIT Press.
- Simpson, G. B., & Kang, H. (1994). Inhibitory processes in the recognition of homograph meanings. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory mechanisms in attention, memory and language* (pp. 359-382). San Diego, CA: Academic Press.
- Slamecka, N. J. (1975). Intralist cueing of recognition. *Journal of Verbal Learning and Verbal Behavior*, 14, 630-637.
- Sloman, S. A., & Rumelhart, D. E. (1992). Reducing interference in distributed memories through episodic gating. In A. Healy, S. Kosslyn, & R. M. Shiffrin (Eds.), *From learning theory to connectionist theory: Essays in honor of William K. Estes* (pp. 227-248). Hillsdale, NJ: Erlbaum.
- Smith, A. D. (1971). Output interference and organized recall from long-term memory. *Journal of Verbal Learning and Verbal Behavior*, 10, 400-408.
- Smith, A. D. (1973). Input order and output interference in organized recall. *Journal of Experimental Psychology*, 100, 147-150.
- Smith, A. D., D'Agostino, P. R., & Reid, L. (1970). Output interference in long-term memory. *Canadian Journal of Psychology*, 24, 85-87.
- Spellman, B. A., & Holyoak, K. J. (1993). An inhibitory mechanism for goal-directed analogical mapping. In *Proceedings of the Fifteenth Annual Meeting of the Cognitive Science Society* (pp. 947-952). Hillsdale, NJ: Erlbaum.
- Stuss, D. T., & Benson, M. D. (1986). *The frontal lobes*. New York: Raven Press.
- Tipper, S. P. (1985). The negative priming effect: Inhibitory priming and ignored objects. *Quarterly Journal of Experimental Psychology*, 37A, 571-590.
- Tipper, S. P., Baylis, G., & Gordon, C. (1987). Individual differences in selective attention: The relation of priming and interference to cognitive failures. *Personality and Individual Differences*, 8, 667-675.
- Tipper, S. P., & Driver, J. (1988). Negative priming between pictures and words in a selective attention task: Evidence for semantic processing of ignored stimuli. *Memory & Cognition*, 16, 64-70.
- Tipper, S. P., MacQueen, G. M., & Brehaut, J. C. (1988). Negative priming between response modalities: Evidence for the central locus of inhibition in selective attention. *Perception and Psychophysics*, 43, 45-52.
- Tipper, S. P., Weaver, B., Cameron, S., Brehaut, J. C., & Bastedo, J. (1991). Inhibitory mechanisms of attention in identification and localization tasks: Time course and disruption. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 681-692.
- Todres, A. K., & Watkins, M. J. (1981). A part-set cueing effect in recognition memory. *Journal of Experimental Psychology: Human Learning and Memory*, 7, 91-99.
- Tulving, E. (1974). Cue-dependent forgetting. *American Scientist*, 62, 74-82.
- Tulving, E. (1989). Memory: Performance, knowledge, and experience. *European Journal of Cognitive Psychology*, 1, 3-26.
- Tulving, E., & Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review*, 80, 352-373.
- Walley, R. E., & Weiden, T. D. (1973). Lateral inhibition and cognitive masking: A neuropsychological theory of attention. *Psychological Review*, 80, 284-302.
- Watkins, M. J. (1975). Inhibition in recall with extralist "cues." *Journal of Verbal Learning and Verbal Behavior*, 14, 294-303.
- Watkins, M. J. (1978). Engrams as cuegrams and forgetting as cue-overload: A cueing approach to the structure of memory. In C. R. Puff (Ed.), *The structure of memory* (pp. 347-372). New York: Academic Press.
- Wickens, C. D. (1984). Processing resources in attention. In R. Parasuraman & D. R. Davies (Eds.), *Varieties of attention* (pp. 239-257). New York: Academic Press.
- Woodworth, R. S. (1938). *Experimental psychology*. New York: Holt.
- Yee, P. L. (1991). Semantic inhibition of ignored words during a figure classification task. *Quarterly Journal of Experimental Psychology*, 43A, 127-153.

(Appendixes follow on next page)

Appendix A

Stimulus Items for Experiment 1

Item type	Category					
	Red	Food	Fly	Animal	Loud	Tool
Dissimilar	blood sunburn fire	bread crackers peas	frisbee glider kite	giraffe hamster sheep	thunder traffic yell	file pliers screwdriver
Similar	apple cherry tomato ketchup radish strawberry		bat pigeon wasp butterfly eagle ladybug		chainsaw drill jackhammer compressor lawnmower sandblaster	

Appendix B

Stimulus Items for Experiments 2-4

Item type	Category					
	Cotton	Leather	Green	Soups	Loud	Sharp
Dissimilar (Experiments 2, 3a, 3b, 4)	curtains napkin sheet	briefcase saddle whip	dollar emerald lawn	chicken clam turkey	jackhammer siren traffic	needle tack thorn
Similar (Experiments 2, 3a, 4)	(clothing*) pajamas robe slacks	belt boots skirt	(vegetable*) artichoke lettuce pepper	mushroom onion tomato	(weapon*) cannon grenade gun	dagger spear sword
Replacement (Experiment 3b only)	pillowcase rag rope	football suitcase wallet	alligator frog martian	minestrone noodle wonton	firecracker whistle yell	pin scissors splinter

* Shared implicit category.

Received December 6, 1993
 Revision received May 31, 1994
 Accepted June 3, 1994 ■