13 Beyond Associations: Strategic Components in Memory Retrieval

Reder, L.M. (1987). Beyond Associations: Strategic components in memory retrieval. In D. Gorfein & R. Hoffman (Eds.), *Memory and learning: The Ebbinghaus Centennial Conference,* Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 203-220.

Lynne M. Reder
Carnegie-Mellon University

There are two strong traditions within the contemporary field of memory research, one belonging to Bartlett (1932) and one to Ebbinghaus (original published in 1885; English translation in 1964). The Ebbinghaus tradition is best characterized by its rigor, its precision, and its attempts to uncover invariants in the memory system. His development and use of the "nonsense syllable" reflect his attempts to control for the influence of prior knowledge on list learning. The Bartlett tradition, started a half-century after Ebbinghaus, is quite complementary in its perspective. There is less focus on precise memory "laws," and more concern with how knowledge interacts with learning and memory.

An important component to Ebbinghaus' theory of memory is the formation of associations, direct and remote; learning is essentially the formation of associations. Bartlett's position is that we do not store verbatim memories. Our memories are *reconstructed* on the basis of what we already know. Both theories have had enormous impact on current views of memory, even though their perspectives are largely construed as antithetical.

The position I argue for here is that at times Bartlett's view is correct and at times Ebbinghaus' view is correct. Sometimes memory data seem best explained by Bartlettian principles and sometimes by those described by Ebbinghaus, because people can adopt one of several memory strategies. For example, people can try to retrieve information in a precise search of memory, looking for exact facts studied, or they can try to reconstruct what they have learned in an imprecise way, making use of prior knowledge, as Bartlett suggested. This chapter will illustrate that the same knowledge structure can produce dramatically different results, depending on the processes (or strategies) that are operating on that structure.

The notion that we have multiple strategies to retrieve information from memory has been neglected to a large extent. Only recently has the availability of multiple strategies been a topic of discussion. The contexts under which one strategy is preferred to another has been explored even less. Many theorists who discuss the availability of multiple strategies have assumed that the order of application of strategies is fixed or invariant. The case will be made that not only are there multiple strategies for retrieving the same information, but that the order of selection of strategies is variable. Nonetheless, the factors that contribute to preference for one strategy over another can be understood.

THE ROLE OF ASSOCIATIONS IN MEMORY RETRIEVAL

Numerous scholars have followed Ebbinghaus' tradition and have documented the importance of associations for understanding learning and memory phenomena. One domain within this broad category is that of interference. Much forgetting or difficulty in learning can be thought of in terms of competing responses to the same stimulus association (see Crowder, 1976, for a review). The original paired-associate learning paradigms that demonstrated interference tended to measure performance as probability of recall of nonsense syllables; however, more recently, analogous results using response times have also been found for recognition of sentences (e.g., Anderson, 1974, 1976; Lewis & Anderson, 1976; Thorndyke & Bower, 1974). The paired-associate research showed that the probability of recalling a response to a cue declined if prior or subsequent associations were also learned to that cue or stimulus. The reaction time research showed that the more facts committed to memory about a particular concept, the slower a person is to *recognize* or reject (as not studied) any statement sharing that concept. This result was dubbed the "fan effect" by Anderson because of the assumed underlying propositional representation in which facts are stored as a set of links between concepts, and facts that share the same concepts all "fan out" from the concept node. This finding seemed quite robust. The monotonically increasing RT function with increasing fan was obtained not only for facts about fantasy characters, it was obtained for real facts about famous people (Lewis & Anderson, 1976; Peterson & Potts, 1982). That is, it takes a subject longer to verify that *George Washington chopped down a cherry tree* the more fantasy facts that were also studied about George Washington.

The theoretical explanation for the fan effect is as follows. Information is retrieved by spreading activation from concepts in working memory through the network of associated facts. The time required to retrieve information is a function of the level of activation that the concept nodes receive. Fanning of multiple paths from a concept node dissipates the activation that the node sends down any one path and increases retrieval time.

The fan effect would never have been challenged if everyone subscribed to the Ebbinghaus tradition of using nonsense syllables or at least materials devoid of any inherent interest. Anderson's original materials were sensible statements, but they were random combinations (generated by a computer) of nouns and verbs (screened for sensibility by a human). Smith, Adams, and Schorr (1978) were intrigued by the paradoxical implications of the fan effect, that knowing more was detrimental. They replicated Anderson's findings, but only with statements that were thematically unrelated to one another. Subjects in one condition studied pairs of facts about fictitious individuals, such as: *Marty did not delay the trip* and *Marty broke the bottle*. In another condition, subjects studied these two facts plus a third unrelated fact such as: *Marty painted the old barn*.

Replicating past fan results, subjects were slower in the three-fan condition than in the two-fan condition. In another condition, however, subjects studied a fact that integrated the first two facts into a theme, such as: *Marty christened the ship*.

When the third fact was thematically related to the other two, there was no difference in verification times between the two-fan condition and the three-fan condition.[1] Moeser (1979), too, found that thematicity changed the effects of multiple associations.

The interpretation of the results offered by Smith et al. seemed to follow the Bartlett tradition for the thematically related materials. They adopted the model proposed by Anderson for random pairings, but assumed an entirely different schema-like representation and process for the thematically related or more meaningful materials. Although Smith et al. partly resolved the "paradox of the expert" by showing that knowing more did not interfere when the material to learn was "integrated," several questions remained unanswered. Why would we have completely different knowledge representations and processes depending on whether the material is thematically integrated or not? On the other hand, why were people not *better* at making judgments when they had learned more integrated material than when they had learned less? Thematically related material only made the interference effects smaller. The next section will answer these questions.

Different Strategies for Responding

It is possible to explain the results found by Smith et al. without assuming different knowledge representations in the two conditions. Other work of my own (Reder, 1976, 1979) led me to believe that people may use inferential reasoning to answer a question even when the information is directly stored. The

[1] In the Lewis and Anderson study, the fantasy facts about famous characters such as George Washington did not form a consistent theme, i.e., they were random combinations of other predicates.

theory that people often prefer to make plausibility judgments rather than to search for a specific fact helps to explain the "paradox of the expert." The Smith et al. result that the fan effect attenuates with thematically related facts can be explained by assuming that subjects often adopt a plausibility strategy to recognize the facts rather than actually search for a specific fact. That is, subjects decide that "it is *plausible* that I studied this fact if it is thematically consistent with other facts I know I studied." If finding any fact in memory about Marty that is consistent with ship-christening would suffice to "recognize" a specific Marty-ship-christening statement, then, of course, there would be little effect of the number of ship-christening facts associated with Marty. Thus, my analysis of the Smith et al. result is that subjects were using different strategies in the two conditions, not different knowledge representations.

This point of view is supported by the results of Reder and Anderson (1980). We replicated the results of Smith et al. using thematically related materials. However, we only replicated their results when the not-studied test foils were unrelated to the studied theme. A subject might study the following three facts: (1) *The teacher went to the train station.* (2) *The teacher bought a ticket for the 10:00 train.* and (3) *The teacher arrived on time at Grand Central Station.* In the block of trials where the studied statements were tested with unrelated foils, a foil to be rejected (as unstudied) might be: *The teacher called to have a phone installed.*

An unrelated foil like this would allow the subject to judge plausibility or thematic consistency instead of truly making a "recognition" judgment. When the foils were thematically related to the studied theme, such as *The teacher checked the Amtrak schedule.* subjects could not use thematic consistency or plausibility to make accurate recognition judgments. (In all conditions, the to-be-rejected foils were constructed by re-pairing studied predicate and occupation terms. In this way, foils could not be rejected because the subject knew a specific word or phrase was not used during the study.) When the foils precluded use of a plausibility strategy, the fan effect was as large with thematically related material as with unrelated material. Fig. 13.1 shows the different RT functions for recognition judgments depending on the type of foil that was tested with the studied facts.

Very different fan functions were obtained with the same materials in blocks of trials that differed only in terms of the type of not studied foils to reject. Since these different functions were produced by the same subjects, an explanation based on different long-term memory representations for thematically related materials is unlikely. Instead, it seems more reasonable to propose that interference or fan effects obtain in recognition tasks if the materials preclude the use of a plausibility strategy.

The question of why knowing more did not *facilitate* still remained. Other work (Reder, 1982) suggested a possible explanation: Strategy-selection is affected by a number of variables, in addition to the type of foils used. For

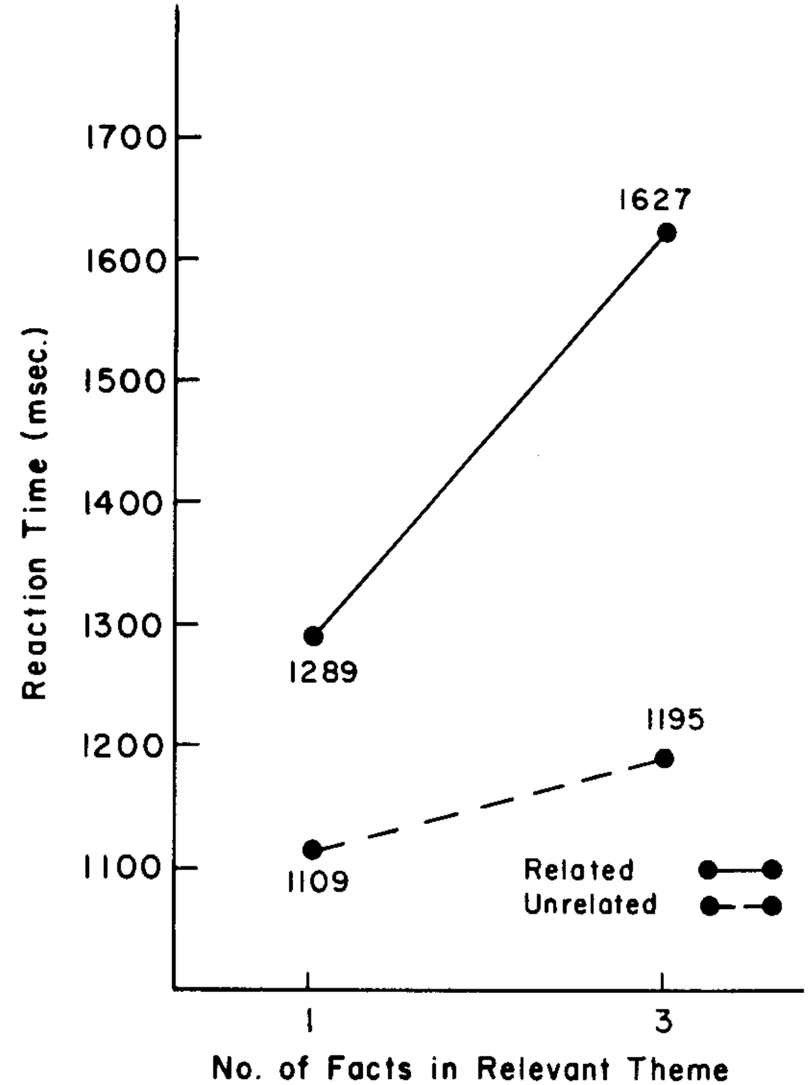

FIGURE 13.1. Mean RT for correct recognition judgments, plotted as a function of the number of predicates associated with the probed character, and whether the block of trials used thematically related or unrelated foils. The data are averaged over yes and no responses. (Adapted from Reder & Anderson, 1980, Experiment 2, Fig. 4.)

example, the task demands (such as recognition versus plausibility judgments) or the strength of the relevant memory traces can affect which strategy is selected. In many situations there is a *mixture* of strategy use, such that sometimes one strategy is selected and sometimes the other is selected.[2] The flat fan function with thematically related materials might have resulted from a mixture of using the recognition (direct retrieval) strategy some of the time and the plausibility strategy the rest of the time. The fan function when plausibility is precluded shows a positive slope, so perhaps the fan function when plausibility is used exclusively would show a negative slope (yielding a roughly flat function when averaged together).

This speculation, that plausibility judgments would actually be facilitated by knowing more, was tested by Reder and Ross (1983). In this study, too, all subjects learned thematically related sets of information and were tested with those facts in a variety of conditions. As in Reder and Anderson, for some blocks of trials subjects were required to make recognition judgments in the presence of plausible (thematically related) foils, and for other blocks of trials they made recognition judgments in the presence of implausible (thematically inconsistent) foils. In addition, a new condition was used in which subjects were actually told to make thematic consistency judgments rather than recognition judgments. They were to say "yes" to both the studied statements and the plausibly true, namely the thematically related but unstudied statements. They were to say "no" to thematically unrelated statements. Figure 13.2 plots the fan functions for the three types of statements (studied, related, unrelated) used in the thematic relatedness block of trials. (The other blocks of trials replicated the earlier studies.) Here, the fan function for the thematically related not-studied statements showed a sharp, *negative* slope. Those statements could only be accepted by a plausibility-like (or thematic consistency) strategy. So the hypothesis that a plausibility strategy would show facilitation with fan was confirmed.

It is worth noting that the slope for stated probes was also negative, but much less steep than for related, not-studied items. This too can be accounted for by assuming that subjects used a *mixture* of the two strategies, since either one produces a correct response for studied statements. The bias to use plausibility was greater in the blocks where subjects were actually asked to judge thematic relatedness. To the extent that subjects were biased to use the plausibility strategy more often as a first strategy in the block requiring those judgments, the slope for the stated probes should be more negative than in the recognition block. Since only the plausibility strategy produces a correct response for the plausible, not-studied items, the function is much more steeply negative for them. Response times for these statements are also much slower than for the other test items because two strategies must often be tried before a correct response is given. That is, first the direct retrieval strategy is tried, but the statement is not

[2]Some of the data supporting these claims will be described later.

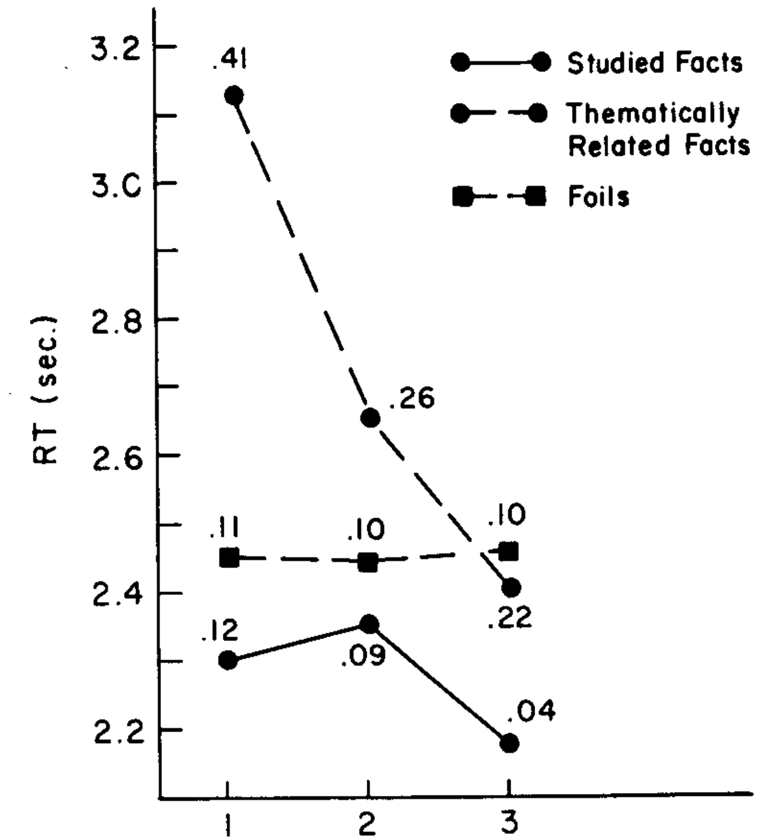

FIGURE 13.2. Mean RTs (and proportion of errors) as a function of relevant fan and type of probe in the consistency task. (Adapted from Reder & Ross, 1983, Fig. 3a.)

found in memory. If the subject quits with a guess of "no," an error is recorded and the time is not averaged into the mean RT. If the subject elects to go on to make a plausibility judgement, the RTs are necessarily much longer than when only one strategy is used to make a decision. Note the high error rate for these statements as well.

Strategy-Selection is Variable

The account given above implies that strategy-selection for question-answering does not proceed in a fixed order. Previous conceptualizations have assumed that people first attempt to answer a question by searching directly for the statement in memory. Only when it is clear that the fact cannot be found is an inferential strategy evoked (e.g., Lachman & Lachman, 1980, Lehnert, 1977). There are a number of results that support the notion that direct-retrieval is not always the first strategy of choice. For example, data of Reder (1979) indicated that subjects make inferences even when the information is stored in memory. In those studies, subjects were asked to read short stories and make judgments about the plausibility of assertions on the basis of the stories that they read. Some of the statements to be judged had actually been presented in the story as part of the story (randomly determined for each subject). The plausibility of the test sentence with respect to the story affected judgment time even when the item had been explicitly presented. Although there was a clear RT advantage for stated (explicit) probes over not-stated probes, the plausibility of the statements affected stated probes as well. Figure 13.3 plots the data from the stated and not-stated conditions for the highly and moderately plausible probes when tested immediately after reading the relevant story.[3]

One explanation for faster RTs for highly plausible statements assumes that the probability of drawing the inference and then finding it in memory is greater for the highly plausible statements, and that subjects always try to search memory for a specific fact first. The problem with this explanation is that it predicts no plausibility effect for probes that had been stated in the story. A different explanation for both the plausibility effect and the speed advantage for stated probes involves a simple race between the direct retrieval process and the plausibility process, where both processes execute in parallel. By assuming that sometimes one process wins and sometimes the other wins, both effects can be accounted for, the faster times for presented statements and for highly plausible statements.

The simple parallel race model just described can be ruled out, however, if one considers the data of Reder (1982). Those experiments were quite similar to Reder (1979), except that some of the subjects were asked to make recognition judgments instead of plausibility judgments. Judgments were either made right after reading a story, after reading 10 stories, or 2 days after reading all 10 stories. In some conditions, subjects actually were faster at plausibility judgments at longer delays than at shorter delays. This is a result that a simple race model cannot explain. Subjects were faster at a delay in those situations where the direct retrieval strategy could not produce a correct response, namely for not-stated plausible inferences. Figure 13.4 presents the response times and error rates for the two tasks (plausibility and recognition) as a function of delay of test and whether the probes had been stated in the story.

An explanation that can account for this result is to assume that at the shorter delay intervals, subjects are inclined to try the direct retrieval strategy first. That

[3]The experiment also included "primed" inferences, and inferences that were verb-based, i.e., they immediately followed from the verb in an assertion. The statements were also tested at various delays. The subset of data graphed here seemed most representative and relevant.

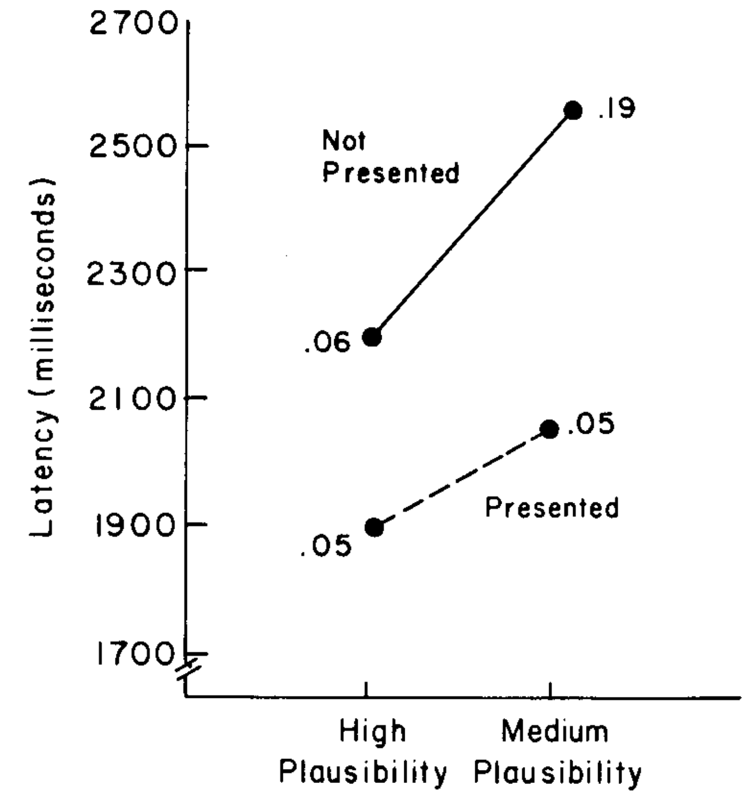

FIGURE 13.3. Mean RT for correct plausibility judgments (and error rates), plotted as a function of plausibility of the test probe and whether it had been presented in the story. (Adapted from Reder, 1979, Experiment 1, Fig. 1.)

strategy will produce the correct response in many conditions. However, when subjects are asked to make plausibility judgments and fail to find the probe in memory, they must go on to try the plausibility strategy or risk making an error. At longer delay intervals, there is an increased tendency to try the plausibility strategy first. This means that for not-stated plausible inferences, the useless direct-retrieval strategy is avoided, making overall response times faster in that condition.

Other aspects of Reder (1982) also supported the notion that subjects became more inclined with delay to adopt the plausibility strategy first in both the plausibility task and the recognition task. As with Reder (1979), the plausibility of the test questions was also varied. The effect of plausibility, (i.e., the difference in response times between the highly and moderately plausible state-

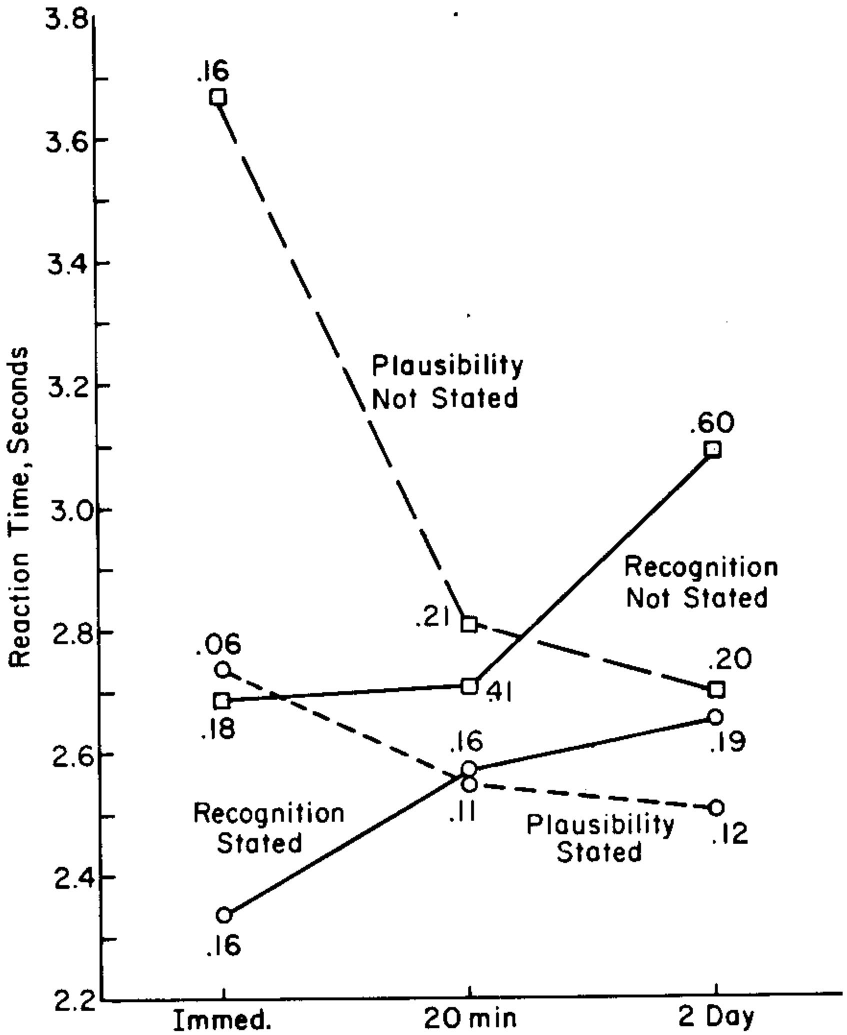

FIGURE 13.4. Mean RT for plausibility and recognition judgments as a function of whether the probe had been stated in the story, plotted across levels of delay. (Adapted from Reder, 1982, Experiment 1, Fig. 3.)

ments) increased at the longer delay intervals. Also, accuracy declined greatly for not-stated plausibles in the recognition task, especially for the highly plausible statements. Figure 13.4 shows that not-stated items in the recognition task were erroneously accepted 60% of the time. The error rates were 50% for the moderately plausible and 70% for the highly plausible.

Given the theoretical interpretation of Reder (1982), that subjects were more inclined to adopt the plausibility strategy as memory traces faded (i.e., at longer delays), it seemed reasonable to predict that at longer delays the pattern found by Reder and Ross (1983) in the fan paradigm would also show a stronger influence of the plausibility strategy (i.e., more of a negative slope). Reder and Wible (1984) conducted an experiment similar to Reder and Ross, in which the major difference was that subjects were tested 48 hours after learning the material as well as being tested on the day of learning.

As expected, the effect of fan (the number of sentences sharing the same concepts) was strongest in those conditions where only one strategy could produce the correct response. Subjects who were asked to make recognition judgments showed the greatest interference from increased fan for the thematically related, not-studied items. This replicated earlier results. Inconsistent items and stated items could be correctly recognized using the plausibility strategy as well as the direct retrieval strategy, so those RT functions were flatter. On the other hand, those subjects who were asked to make thematic consistency judgments showed the most facilitation from increased fan for the related, not-studied items, since only the plausibility-like strategy would work.

The prediction that subjects would be more inclined to use the plausibility strategy at longer delays was supported by a number of results in Reder and Wible (1984). Figure 13.5 shows the mean facilitation (or speedup) from the first session to second session as a function of task and probe type. For all items, subjects are somewhat faster during the second session, possibly due to practice or due to greater fatigue during the first RT session after learning the materials. It is reasonable to consider the flat line for the stated probes as a baseline of no true facilitation or loss due to strategy shifts, because either strategy works equally well for these items. The most interesting changes in relative speed from the first to the second session are for the related, not-stated items. In the recognition task, where the plausibility strategy would produce the wrong response for these items, there is a relative hindrance due to the strategy shift, i.e., much less speedup than the baseline. In the consistency task, where only the plausibility strategy produces the correct response, there is the greatest speedup. Presumably this occurs because there are far fewer trials where subjects first use the inappropriate direct retrieval strategy prior to the plausibility strategy.

The speedup results just described are inconsistent with a simple, parallel race model for the same reasons that the speedup in Reder (1982) is inconsistent with it. Other aspects of the data also argue for a shift in strategy preference (from direct retrieval to plausibility) with increasing delay. The slopes of the fan

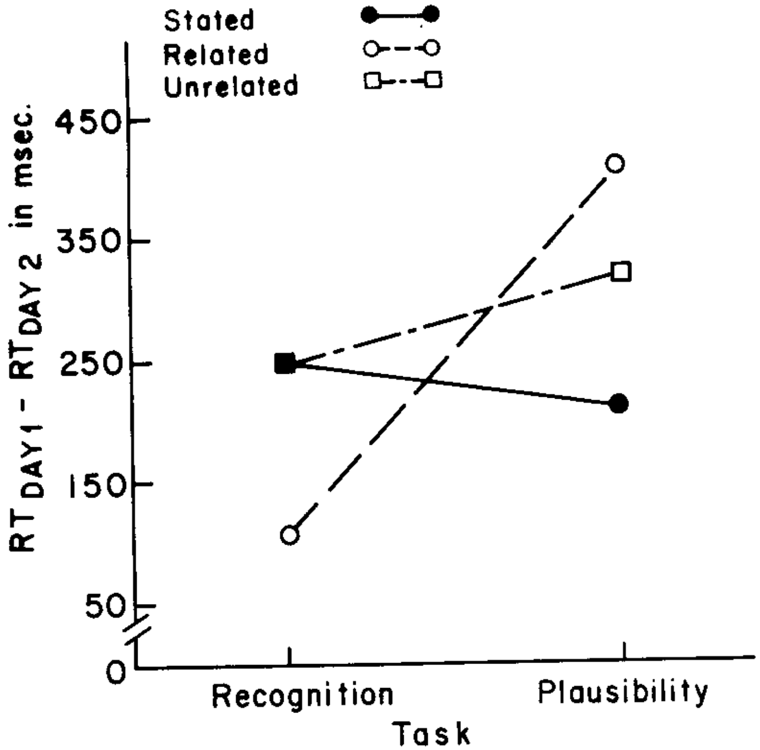

FIGURE 13.5. Mean facilitation or speed-up from first session to second session, plotted as a function of task and probe type. (Adapted from Reder & Wible, 1984, Fig. 2.)

functions, negative for the consistency task and positive for the recognition task, both become more negative at the 2-day delay. This is because with a greater portion of trials using the plausibility/consistency strategy at a delay, there should be more of a facilitation effect from increased fan.[4]

The accuracy data also showed a shift toward greater use of the plausibility strategy at longer delays. Figure 13.6 displays the accuracy data as a function of relevant fan, task, delay and type of test probe. During the first session, in which the subjects tended to use the direct retrieval strategy for the plausibility task, accuracy was poor for thematically related, not-stated probes, since it produced

[4]In the recognition task, the slope changed from a +135 msec. slope to a −65 msec. slope at a delay. In the consistency task, the slope showed less influence of interference initially, starting with a +12 msec. slope and shifting to a greater facilitation than in the recognition task, a −146 msec. slope. These slopes are only computed for the stated and inconsistent probes since only one strategy can be correctly used for the not-stated inconsistent probes.

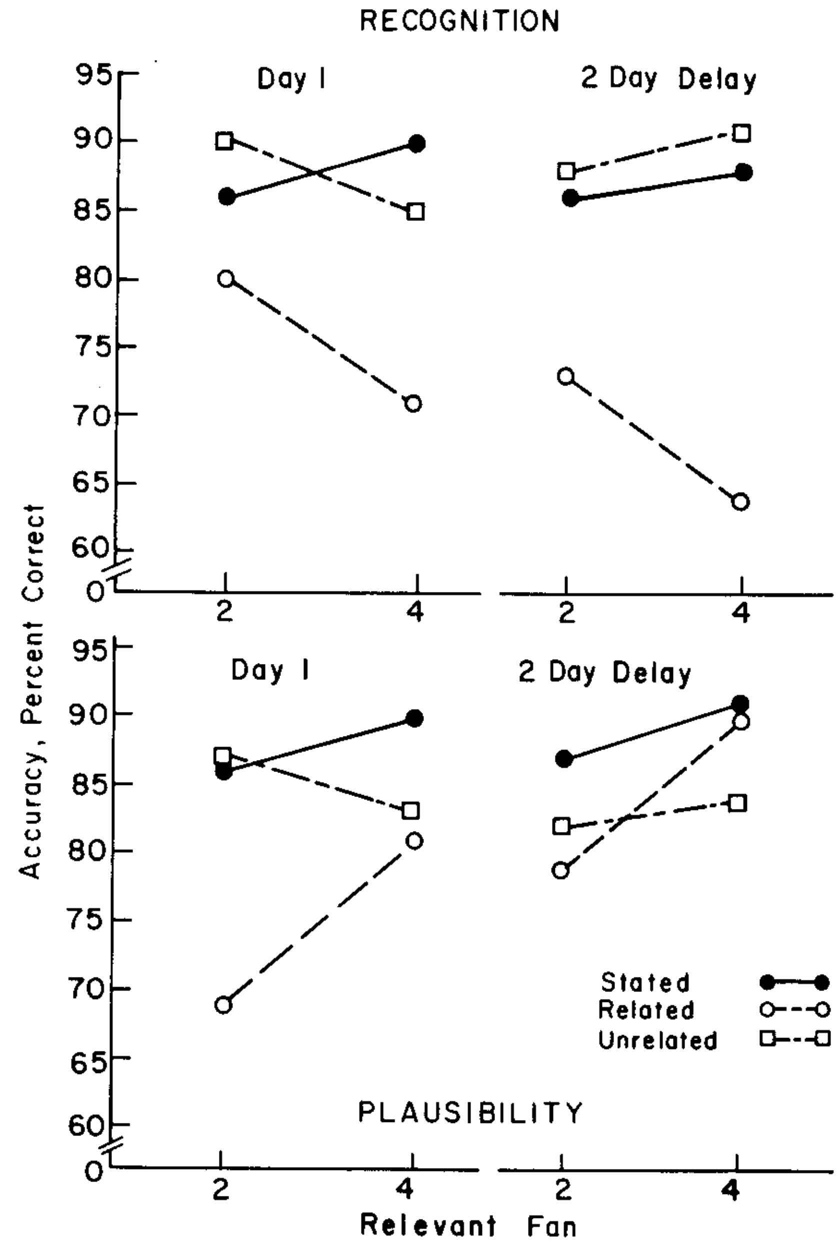

FIGURE 13.6. Mean percentage correct for judgments in the recognition task (top) and consistency task (bottom), plotted as a function of relevant fan for each probe type. The short-delay data are displayed in the left quadrants, and the long-delay test data are plotted on the right. (Adapted from Reder & Wible, 1984, Fig. 3.)

the wrong responses (no's instead of yes's). Accuracy improved 10% for both levels of relevant fan with delay for those statements where both strategies could not work.

The accuracy pattern in the recognition task showed the opposite trend for the related statements for exactly the same reason. These were the only probes that would produce an error when the plausibility strategy was invoked. At the delayed test, there was a greater tendency to use the plausibility strategy and so more errors were made. Furthermore, these probes were 10% worse in accuracy for the high-fan condition, that is, when there was more information consistent with the theme of the not-studied probe.

THE MECHANISMS FOR STRATEGY-SELECTION

Several conclusions seem clear given the results just described. Questions are not always answered by using the same question-answering strategy or process. Which strategy is used to answer a given question varies. The available strategies do not compete for execution by racing against each other to see which one will complete first. A nonfixed order serial model can easily account for the data described. A parallel race model can also account for the data if we assume that the *allocation of processing resources* is unequally distributed for the two strategies. By assuming a shift in the allocation of resources from one strategy to another instead of a shift in which strategy is tried first, the same pattern of speedups, and so forth, can be explained.

Regardless of whether or not one assumes a variable, serial strategy-selection model or a differential allocation of resources in a parallel-race model, it still follows that there must be a preliminary process that determines strategy selection or allocation of resources among strategies. This preliminary process has two sub-components or stages: an initial evaluation of knowledge relevant to the question followed by a decision of which strategy to follow. I propose that the initial evaluation is an automated process while the decision is a conscious, controlled process (e.g., Neely, 1977; Posner & Synder, 1975; Shiffrin & Schneider, 1977).

Initial Evaluation

A number of factors are involved in the initial evaluation. One involves assessing the familiarity of the words in the question. The more familiar the words, the more a person is biased toward using direct retrieval. The initial evaluation also involves assessing how many intersections in memory there are among the words from the question. The more intersections, the more the person is biased toward plausibility.

The idea that we can automatically determine the familiarity of the concepts provided in the question has been proposed elsewhere (e.g., Hasher & Zacks, 1979; Jacoby & Dallas, 1981; Mandler, 1980). The activation level of the terms in memory that were referred to in the probe can be compared with their "resting" activation levels. If they seem higher than expected at the time of questioning, it is assumed that the words were encountered recently.

The proposal that relatedness affects decision times is also not new (e.g., Rips, Shoben & Smith 1973). In my view, the "relatedness" of the concepts in the question is monitored through the interconnections in memory. Relatedness is defined as the degree to which words in a question cause *activation* to intersect in memory. The more intersections that are detected in memory as a result of a query, the more potentially relevant information there is available for question answering.

Familiarity detection and intersection detection are processes that monitor the automatic spread of activation from the concepts in the question. This spread of activation is assumed to be automatic, as are processes that monitor the level of activation and the extent of intersections. The bias to use the direct retrieval strategy "trumps" the plausibility strategy since direct retrieval is a faster and easier strategy than judging plausibility when the queried fact is relatively accessible. This is because when memory search is relatively easy, the plausibility strategy does not have the search-time advantage to counteract its long plausibility computation time.

Strategy Selection

In deciding which strategy to apply, the subject integrates the biases from the initial evaluation along with considerations or factors that are *extrinsic* to the test question. These extrinsic factors include things such as task instructions and probability that a particular strategy will be successful. Some of these variables have already been shown to influence strategy selection, for example, form of the instructions (Gould & Stephenson, 1967; Reder, 1982), ease of discrimination among alternatives (Lorch, 1981; Reder & Ross, 1983; Reder & Wible, 1984), impressions of one's own expertise (Gentner & Collins, 1981), and form of the question (Rips, 1975).

In addition to these variables, it seems reasonable that strategy selection would be affected by recent prior history of success with a strategy, nominal constraints of the task, special knowledge that a strategy will or will not work, and motivation to perform well. The influence of extrinsic factors on strategy selection is partly a function of how strong the bias is from the automatic assessment of "feeling of knowing" from the first stage and how compelling the factors are from this stage. If there is overwhelming evidence that a strategy will not work, or if subjects are heavily penalized for making errors, they may ignore the biasing information from the automatic assessment.

evaluation with facts that we can more consciously assess in order to select one strategy to try first (or to which to devote most of our processing capacity).

GENERAL CONCLUSIONS

This volume is in honor of the hundredth anniversary of the publishing of Ebbinghaus' famous treatise on human memory. The message of this contribution is that we should reconsider Ebbinghaus' ''laws'' from the perspective that memory performance reflects cognitive strategies and that these strategies vary from situation to situation. For example, the same structure of associations can either hurt memory performance or facilitate it, depending on other constraints of the task. It is no wonder then that the conclusions of Ebbinghaus and those of Bartlett seem so contradictory. Each had constructed tasks and tested subjects in situations that encouraged completely different strategies. Ebbinghaus constructed tasks that required veridical, verbatim recall that minimized the usefulness of prior knowledge. Bartlett used tasks that encouraged reconstruction of the information by using prior world knowledge. Both sets of ''laws'' or principles are useful, so long as we acknowledge that they apply only in contexts that encourage the corresponding memory strategy.