Journal of Experimental Psychology:
Learning, Memory, and Cognition
1983, Vol. 9, No. 1, 55–72

0278-7393/83/0901-0055$00.75

Integrated Knowledge in Different Tasks: The Role of Retrieval Strategy on Fan Effects

Lynne M. Reder
Carnegie-Mellon University

Brian H. Ross
Stanford University

Prior research has shown that the *fan* effect (slower response times [RTs] to verify facts when more are studied on a topic) is attenuated when thematically related facts are used. This alleviation of interference occurs only when subjects can use consistency judgments instead of direct retrieval to make recognition judgments. Exploring this issue further, we discovered that knowing additional facts relevant to the test fact can actually speed judgment times when subjects are asked whether the test fact is consistent with what is known rather than if it had been studied. The same subjects displayed three different RT functions for the same memory items when they performed in three different test blocks intended to invoke three different strategies. We also varied the number of topics studied about fictitious characters while holding constant the number of total facts. Unlike simple fan, we found that the greater the number of topics studied with a character, the longer subjects took to respond, regardless of the strategy used in that test block. We present a model that is in accord with this pattern of data.

The mechanisms involved in storing and retrieving information from memory are sensitive to the relations between the contents of memory. For example, subjects are slower to recognize a fact when other facts share some of the same concepts (e.g., Anderson, 1974, 1976; Anderson & Bower, 1973; Hayes-Roth, 1977; King & Anderson, 1976; Lewis & Anderson, 1976; Thorndyke & Bower, 1974). This effect has been called the *fan effect* because of the underlying representation that has been assumed to explain this phenomenon (Anderson, 1974). This representation involves a network of propositions (facts), where nodes in the network represent concepts, and links represent relations among concepts. Facts that share the same concepts have relational links fanning out of the same concept nodes. Activation spreads from various concepts through the network along the links or pathways that encode propositions. Activation spreads rapidly through the network; the amount of activation that passes down any one link depends on the number of other links that share the activation from the concept. The level of activation of a particular knowledge structure determines the rate at which it is processed and the probability of successful processing (see Anderson, in press).[1]

This fan effect has been under close scrutiny of late (e.g., McCloskey & Bigler, 1980; Moeser, 1977, 1979; Smith, Adams, & Schorr, 1978) because it seems paradoxical. The more one knows about a specific topic, the faster one should be able to answer questions about that topic. Yet according to the theory, the more one knows about a particular topic, the slower the retrieval of any given fact.

Smith et al. (1978) and Moeser (1977, 1979) showed that when the facts associated with a particular concept are thematically related, there does not seem to be any interference between them. That is, previous re-

Parts of these results were presented at the Cognitive Science Society meeting at Berkeley, 1981.

A National Institute of Mental Health traineeship supported the second author, who is now at the University of Illinois at Urbana-Champaign.

We thank J. R. Anderson, W. Jones, G. L. Murphy, C. Wible, and especially D. A. Norman for helpful suggestions to improve this article.

Requests for reprints should be sent to Lynne Reder, Department of Psychology, Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213.

[1] The prior model of Anderson (1976) assumed that time to retrieve a fact depended on time to activate the set of propositional links and that activation spread at different rates as a function of fan. This view has been revised in light of recent empirical results (e.g., Ratcliff & McKoon, 1981; Schustack, 1981).

search on the fan effect has shown that the more facts studied about a fictitious individual (e.g., Marty), the longer subjects take to judge that a given fact had been studied. However, if all facts associated with Marty could be integrated into a theme (e.g., ship christening), the fan effect is diminished. These results had been interpreted by Smith et al. as supporting a different type of representation for thematic material, specifically a "script" structure for related materials. In that conception, time to access a fact depends on time to access the relevant script associated with a character. The number of studied facts related to the theme does not affect time to find the correct one in the script.

Reder and Anderson (1980) showed that the fan effect diminishes only when subjects treat the recognition task as a plausibility judgment task, rather than retrieving the specific fact. "Plausibility" means judging whether the fact was studied by seeing if it is *consistent* with information known to be studied. Treating a recognition task as a plausibility task will only work when the foil sentences are unrelated to the studied sentence; when the foils are thematically related to the facts studied about the individual, subjects must retrieve particular propositions to respond accurately. With related foils, the fan effect is as large with thematically integrated materials as with nonintegrated materials.

Because the fan effect can be obtained with integrated materials, the script hypothesis put forward by Smith et al. (1978) must be rejected. Instead, Reder and Anderson assumed the same type of propositional network representation used to explain the fan effect with unrelated materials. The only modification to that original representation was thematically related propositions are grouped together into subnodes, and the subnodes are attached to the "individual" nodes.

Subnodes

Figure 1 sketches the type of memory representation proposed by Reder and Anderson to account for their data. This semantic network encodes three facts about Marty related to ship christening and two about washing clothes. Each thematically related set of facts has its own subnode. For different tasks, the same representation and activation process are used, but different strategies (or evaluation criteria) are used. When subjects are not forced to retrieve a specific fact, because the foils are not thematically related to the facts studied about the character, subjects can stop search at the appropriate thematic subnode (e.g., ship christening). The interference due to multiple related facts is avoided because activation need not go beyond the subnode, and the fan of related facts is off of the subnode. Even when subjects must retrieve a specific fact, response times (RTs) are faster than with unrelated material (McCloskey & Bigler, 1980) because the search of thematically inappropriate material is stopped at the inappropriate subnode; that is, irrelevant facts are not inspected, and the activation is refocused to the relevant node.

This subnode model accounts for another Reder and Anderson (1980) finding: The in-

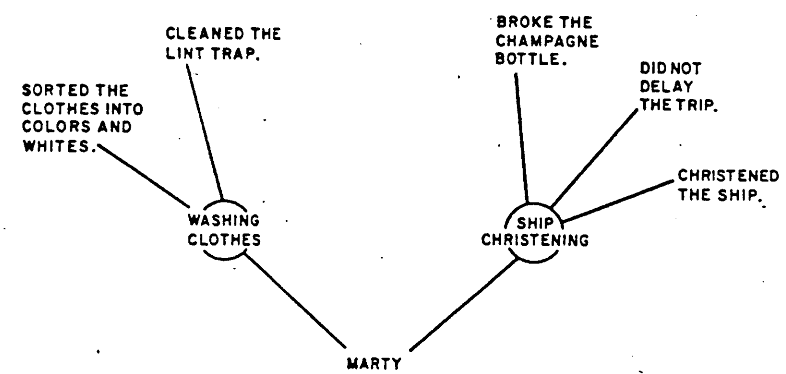

Figure 1. Type of memory representation proposed by Reder and Anderson (1980) that includes thematic subnodes to group-related sets of facts about Marty.

terference effects did not depend on the number of facts learned about a theme unrelated to the one specifically queried. This model also explains the similar result of McCloskey and Bigler. They had one or two sets of facts studied with a particular concept; they obtained a fan effect for the number of facts relevant to the test probe (because their foils were thematically related to the facts being probed) but not for the number of irrelevant facts.

One other related result found in both the Reder and Anderson and the McCloskey and Bigler studies is also accounted for by this subnode model. In Reder and Anderson's study, subjects were slower to verify a fact, the more *themes* there were associated with the fictitious individual. We call this result the *theme-fan effect.* This effect occurs because activation from the person node must be shared by the links fanning out to the theme subnodes. The data of McCloskey and Bigler are also consistent with the model because there was a difference in RT between no irrelevant facts and some irrelevant facts (presence or absence of an extra subnode link).

Strategies

The reduction in the fan effect with thematically related facts can thus be thought of as a change in strategy in which subjects opt to make consistency judgments rather than to retrieve a specific fact. There is other support for the notion that people sometimes answer questions by judging consistency or plausibility even when asked to make recognition judgments (Reder, 1982). Moreover, it seems reasonable that in real situations, people might make plausibility judgments rather than decide whether a specific fact has been presented. Psychological experiments are probably one of the few situations that require people to discriminate what was said from a correct inference or paraphrase of something asserted.

In all fan experiments, subjects have been asked to make recognition judgments, not to judge consistency. The experiments of Reder (1982) showed that factors such as official task as well as "convenience" of a strategy affect the strategy selected. Therefore, the attenuated fan effect of Reder and Anderson (1980) in conditions where subjects can "get away" with using the consistency strategy may reflect a mixture over trials of direct retrieval and consistency. Presumably, if subjects were asked to judge whether a fact was consistent with what had been studied (rather than whether it had been studied), there would be a larger proportion of the trials devoted to the consistency strategy. Because the positive slope for fan decreased when subjects were "allowed" to use the consistency strategy, it seems reasonable that the function might become *negative* when consistency is the official task. That is, knowing more about a topic might actually *facilitate* the speed with which one decides that a related fact is true. This was one prediction of the present experiment.

Strategy Selection and Subnode Structures

The subnode representation and search process described earlier (see Figure 1) can account for the fan effect or for the attenuation of the fan effect with the same representation by assuming that a subject-controlled strategy affects whether search and processing proceed beyond an intersection of activation at the relevant subnode. There is a second, strong prediction that follows from this conceptualization: Although the size of the fan effect should vary with strategy or task demands, the size of the theme-fan effect should not vary with strategy or task and should always be positive. A reexamination of Figure 1 should make this clear. Regardless of whether or not a person elects to continue processing after activation intersects at the appropriate subnode, the time for activation to reach that subnode will depend on the number of links fanning out of the person node. The number of themes learned with a character, of course, is the variable that affects this theme-fan.

Predictions

In this experiment, we expect to find three different effects of what we call *relevant-fact fan* (or *relevant fan,* to be distinguished from *theme-fan*), depending on the nature of the foils and depending on whether subjects were actually asked to make recognition or consistency judgments: (a) an increasing RT function with increasing fan for recognition tested with thematically related foils; (b) a flat function for recognition tested with unrelated foils; and (c) a decreasing function with increasing fan when statements are to be evaluated for consistency rather than for recognition. In all three test blocks, however, we expect the same, positive relationship between number of themes associated with the character and RT. In no case do we expect total number of facts (ignoring the number of themes) to have any effect. The mechanisms by which consistency judgments allow fan to facilitate RTs, rather than just attenuate the usual fan effects, are described in the General Discussion section.

Method

Procedural Overview

There were three major phases in the experiment: initial presentation, learning, and test. In the initial presentation phase, subjects were shown sets of facts about various characters. All facts about a given character were presented together on a CRT screen for a fixed amount of time for the subject to study. Then all facts about another character were displayed. Facts were organized by theme, with titles of themes presented, as well as the facts.

In the second phase, the learning phase, subjects had to demonstrate that they knew which facts were associated with each character. This was accomplished in four subphases that either required total recall of facts studied with a character or required recall of all facts on a given theme learned with that character. Subphase 1 used a "drop-out" procedure with a criterion of two perfect recalls of all facts associated with a specific theme for a specific character. Subjects would see a character name with a particular theme name and a number next to that theme. That meant subjects were to recall that number of facts on that theme associated with the character. After typing in as many facts as they could recall, the subjects were given feedback indicating the actual sentences studied. Then the subjects would score themselves[2], and the computer would display a new character with one of the themes studied with the character. After subjects could recall twice all the sentences for a given theme for a given individual, that particular character–theme pair was dropped out from further testing in this subphase. After all character–theme pairs reached the double drop-out criterion, subjects went to the second subphase.

In Subphase 2, subjects were given the character's name and a number indicating the total number of sentences associated with that character. For example, if subjects studied two facts with Marty at the beach and three facts about Marty taking a train, then the subjects would see "Marty-5." After the subjects typed in the appropriate number of sentences, the terminal displayed the correct sentences, and the subjects scored themselves. After being tested on all characters, subjects went on to Subphase 3. For all characters not perfectly recalled in Subphase 2, subjects were again tested on character–theme pairs. This subphase was exactly like Subphase 1, except that there was a single drop-out criterion rather than a double drop-out criterion, and subjects were only tested on those characters missed in the preceding subphase. In Subphase 4, the final subphase, subjects were again presented with just the character name and a number, and subjects were required to type in the studied facts. This time the responses were printed out to a file for the experimenter's inspection. If a subject did not correctly recall at least 90% in this phase, that subject's data were excluded from the analyses.

In the critical third phase, the test phase, we collected RTs for subjects to make various judgments about the learned material. These different types of judgments were divided into three test blocks. One test block involved making *recognition judgments* with the foils *thematically related* to the sentences actually studied with that character. A second test block involved recognition judgments with the foils *unrelated* to the material studied with the character. A third block required *consistency judgments* rather than recognition judgments. That is, subjects were not asked to judge whether they had studied a particular sentence but only whether the sentence was consistent (i.e., thematically related) with what they had learned about that character. Because the pattern of predictions depends on different tasks and instructions may be critical, Appendix A contains the complete instructions given during the experiment.

In all three test blocks, there were some special trials that required *theme judgments* such that the screen would display a character name and a theme name (e.g., Marty ship christening) rather than an entire sentence. Subjects would answer *yes* to *Marty ship christening* if they had studied ship-christening facts with Marty. We expected the same pattern for theme judgments as for consistency judgments. These trials served as a check on our interpretation that subjects' responses in the consistency block result from searching only as far as the theme node.

The order of test blocks was randomly determined for each subject, as was the location of the theme judgments. There was the constraint, however, that the two recognition blocks be contiguous so that changing strategies from recognition to consistency judgment would be minimized. Both speed and accuracy were emphasized. Subjects were told to respond as fast as they could while maintaining high accuracy. Accuracy feedback was given after every trial. After every 13 trials, subjects were told the total number of errors they had made on the last 13 trials. Subjects were given results between phases and between the recognition and consistency blocks.

---

[2] Self-scoring, of course, presents the problem of a variable criterion depending on subjects' honesty. They were told to score for "gist" and to count as wrong a fact that did not say essentially the same thing as the answer. It would be ideal to have computer scoring or experimenter monitoring of all responses, but this was not feasible.

Design and Materials

Table 1 illustrates some of the material that a subject might see. All study and test items were statements about a person performing some activity. Test items that had not been studied (foils or thematic statements) were always new combinations of studied character names and studied predicates. In this way, word familiarity could not serve as a decision-making cue.

There are three factors that define the condition within which a particular probe is tested: relevant fan, theme-fan, and irrelevant fan. Relevant fan refers to the number of facts studied with the probed character that are related to the theme of the probe. This number varied from one to four. Theme-fan refers to the number of themes associated with the probe character, and it varied from one to three. Irrelevant fan refers to the total number of irrelevant facts associated with the irrelevant themes; it varied from zero to five. Conditions are described by three digits (see Table 1). The digit on the left refers to the amount of fan in the relevant or probed condition. The other two numbers refer to the amount of fan for each of the other two themes associated with that probed character. So, for example, the 3-2-1 condition means the subject studied three facts relevant to the probe, two facts on one theme unrelated to the probe, and one fact on another theme unrelated to the probe. In this case the theme-fan would be three, but in the 4-0-0, theme-fan would be one.

The test variables are orthogonal to the conditions described above and are displayed in Table 2. All studied facts were tested in all four judgment conditions: in a *recognition* test in the presence of *thematically related* foils, in a *recognition* test with *unrelated foils*, in a *consistency judgment* task, and in a *theme judgment* task. With the exception of the theme judgments, these judgment tasks were divided into blocks. Consider the first example in Table 1, the 3-2-1 condition. The subject has studied three facts about Alan taking a train, two about Alan washing clothes, and one fact about Alan skiing. The test probe concerns the train theme (which is why this condition is 3-2-1—it would be 1-3-2 if the test probe were about skiing). The subject should respond positively to the particular test probe, *Alan bought a ticket for the 10:00 train*, in either of the two recognition blocks and in the consistency block. The probe (foil), *Alan called to have a phone installed*, would be used in a recognition test in which the foils are unrelated to the study material. This probe could also be used in the consistency block as an inconsistent probe. The probe, *Alan watched the approaching train from the platform*, would only be tested as a foil in the related foil block; in the consistency block, it would be tested as a thematic probe.

Note that the last type of probe is the only one to which subjects must respond differently, depending on the task. That is, in the recognition phase, that probe is a difficult foil (with a correct answer of *no*), whereas in the consistency block, the correct answer is *yes*. The test probe, *Alan train*, is a theme judgment, and the subject should respond *yes* to it regardless of whether the trial appeared in a recognition block or in a consistency block. The probe, *Alan circus*, would be a foil for a theme judgment (that theme was studied with Brian), and the subject should respond *no*. Foils for theme judgments, as for other judgments, use new combinations of materials already studied.

Table 1
Examples of Studied Facts

Number of facts in the three themes	Studied facts
3-2-1	Alan bought a ticket for the 10:00 train.
	Alan heard the conductor call, "All aboard."
	Alan arrived on time at Grand Central Station.
	Alan added bleach to the rinse cycle.
	Alan sorted his clothes into colors and whites.
	Alan fell while skiing down the steepest stretch.
1-1-0	Brian watched the freaks in the side show.
	Brian wanted to major in psychology.
4-0-0	Steven called to have a phone installed.
	Steven read and signed the lease.
	Steven unpacked all of his boxes.
	Steven mailed out change of address cards.
3-3-0	James compared five different model cars.
	James paid the car dealer in cash.
	James put the license plates on his car.
	James checked the Amtrack schedule.
	James arrived on time at Grand Central Station.
	James watched the approaching train from the platform.

The design has two further constraints. First, each predicate was studied with two characters to ensure that the activation from the predicate would not be the primary determinant of activation time (see Anderson, 1976, chap. 8). Second, for every predicate studied with a character, we needed a thematically related predicate studied with another character to use as a related foil. In the example from Tables 1 and 2, a predicate studied with James in the train theme would be used to probe Alan's train theme.

Subjects were required to learn 52 sentences about 12 characters. These sentences were about 5 different themes, with 25 character–theme pairings. There were 1,032 test trials in the experiment, representing 17 conditions. Three of these conditions are not described here.[3] The

[3] The omitted three conditions refer to a "strength" manipulation intended to make certain facts stronger in memory. The manipulation yielded results in the expected direction, but very weakly. The manipulation itself was quite weak in that "double-strength" facts were not given twice the exposure. Because the manipulation was weak and the result is weak, for simplicity, we ignore this aspect of the design.

Table 2
Test Questions and Responses

Questions[a]	Recognition: With related foils	Recognition: With unrelated foils	Consistency
Studied fact (Alan bought a ticket for the 10:00 train)	Yes	Yes	Yes
Related statement (Alan watched the approaching train from the platform)	No	—	Yes
Unrelated statement (Alan called to have a phone installed)	—	No	No
Theme judgments			
Alan-train	Yes		
Alan-circus	No		

[a] Examples refer to items studied in Table 1.

remaining 14 conditions can be described with the notation explained above: 4-0-0, 1-1-0, 1-3-0, 2-2-0, 2-4-0, 3-1-0, 3-3-0, 4-2-0, 1-2-1, 1-3-2, 2-1-1, 2-1-3, 2-2-2, and 3-2-1. There were 416 test trials in the consistency block and 416 in the recognition blocks, half in the block with related foils, half with unrelated foils. There were 200 theme judgment trials, half tested in the consistency block, and one quarter in each of the two recognition blocks.

For each subject, the computer program randomly assigned themes and facts within a theme to conditions, characters to themes and conditions, order of presentation of trials, and order of task block (recognition first or last). By randomizing separately for each subject, any effects due to materials would be part of the error term, obviating an analysis of variance (ANOVA) for materials (see Clark, 1973.)

Subjects

Thirty-four subjects were recruited from the general Carnegie-Mellon population. Subjects either received course credit and money or money exclusively ($3 per hour) for an experiment that lasted 4 to 7 hr.

Results and Discussion

One subject was excluded from the analysis because he failed to satisfy the criterion of recalling 90% or more correctly on the last subphase of the learning phase. The RTs for correct responses were truncated to 5 sec (1.5% of the data). After this truncation, analyses used the mean times for all correct responses of a subject in a condition. Occasionally a subject made an error on every trial of a condition so that there was no mean RT value for that subject in that condition. Rather than excluding a value for that subject in that condition, we used the value of 5 sec (the truncation time) to reflect the fact that had that subject made any correct responses, the responses probably would have been long. We also analyzed the data by omitting empty cells, and the overall pattern of RTs remained the same.

The trials in the three test blocks comprise nine different task types; for example, in the consistency task, there were three different types: studied facts, thematic statements, and foils. Each task type was analyzed separately because the variances were very different across types. There were three variables that defined how the studied material was to be tested: number of facts relevant to the test probe (relevant fan), number of themes associated with the probed character (theme-fan), and number of facts associated with the probed character but irrelevant to the probed theme (irrelevant fan). Analyses of variance, individual contrasts, and regression analyses were performed on these data to investigate each of these effects with respect to the various tasks. A two-way ANOVA was performed on the data using two levels of theme-fan (2 vs. 3) and three levels of relevant fan (1 vs. 2 vs. 3). A second analysis examined the effects of irrelevant fan, summing over the fan in either irrelevant theme, which made five levels. Table 3 indicates what conditions were used in each ANOVA.[4] Appendix B lists the

[4] Because the experiment was not a factorial design, not all conditions could be used for a given analysis. The conditions wih 4-fan were not used in the two-way ANOVA because there was not a 3-theme condition with 4-relevant fan. The irrelevant fan analyses used scores cor-

RTs and error rates for each of the 14 fan conditions for each type of item (e.g., target, foil) for each task (e.g., recognition tested with thematically related foils).

Relevant Fan

We predict different patterns for relevant fan effects across tasks. The reason for this prediction is that the strategy selected to answer the question will vary depending on the official task and the characteristics of the discrimination required. In any task, different strategies may be used to answer questions, but we predict that the proportion of one strategy relative to another will vary with task characteristics. We expect to see the most interference due to related facts (relevant fan) in the recognition task where foils should preclude making consistency judgments and the least interference in the tasks that explicitly ask for consistency judgments. We use these predictions to present the results in order to make a complex pattern interpretable.

Recognition task. Figure 2 displays the time to make recognition judgments as a function of relevant fan. The error rate (in proportion) is listed above each data point. The left panel plots studied facts and foils tested in the block that used foils thematically related to studied items; the right panel plots studied facts and foils tested in the block containing unrelated foils. The graphs ignore the number of irrelevant facts and number of themes associated with the probed character.[5] Consider first the results in the right-hand panel, recognition in the presence of unrelated foils. There was no effect of relevant fan measured in terms of RT or accuracy. The F statistics for both targets and foils were less than 1.0 with the exception of the target accuracy. In that case, the effect was still not significant, $F(2, 64) = 2.08$, but suggested, if anything, that accuracy was improving with greater fan.

Table 3
Conditions Used in the Analyses of Variance (ANOVAS)

Two-way ANOVA: Theme Fan × Relevant Fan

Levels of theme fan	Levels of relevant fan: 1	2	3
2	1-1-0 1-3-0	2-2-0 2-4-0	3-1-0 3-3-0
3	1-2-1 1-3-2	2-1-1 2-1-3 2-2-2	3-2-1

One-way ANOVA: Effects of irrelevant fan[a]

Levels of irrelevant fan: 1	2	3	4	5
1-1-0 3-1-0	2-2-0 2-1-1 4-2-0	1-3-0 3-3-0 3-2-1 1-2-1	2-4-0 2-1-3 2-2-2	1-3-2

[a] See Footnote 4 concerning correction procedure.

Now consider the results from the left-hand panel, where the foils should preclude using consistency to make recognition judgments. Relevant fan affected RTs of thematically related foils, $F(2, 64) = 5.41, p < .01$, such that subjects took longer to reject a foil about a theme that had more facts studied with it. Subjects also made significantly more errors when trying to reject foils that were strongly thematically related (high relevant fan), $F(2, 64) = 5.59, p < .01$.

Targets displayed a significant effect of relevant fan on accuracy, $F(2, 64) = 3.22, p < .05$, such that subjects became more accurate to accept a target with high thematic relatedness (high relevant fan). The pattern of

rected for effects of theme fan and of relevant fan (due to the nonfactorial nature of the design.) The correction procedure took into account the different weightings of number of themes and different weightings of relevant fan for different levels of irrelevant fan. The correction involved adding or subtracting the appropriate fraction of the size of the effect for the factor in question, and was determined from the means.

[5] The distinction between relevant and irrelevant fan is meaningless when foils are unrelated to themes studied with the probed character. We have plotted the values and analyzed the data using the same conditions that were used in the other tasks, for example, related foils or studied facts. For those conditions with balanced fan, such as 1-1-0 or 3-3-0, assignment to condition is straightforward. For test probes with unbalanced fan, for example, one fact associated with one theme, and three facts associated with a second theme, the computer arbitrarily assigned the item to a possible condition, for example, 1-3-0 or 3-1-0. The pattern of latencies and errors are quite similar when only the balanced fan items are used.

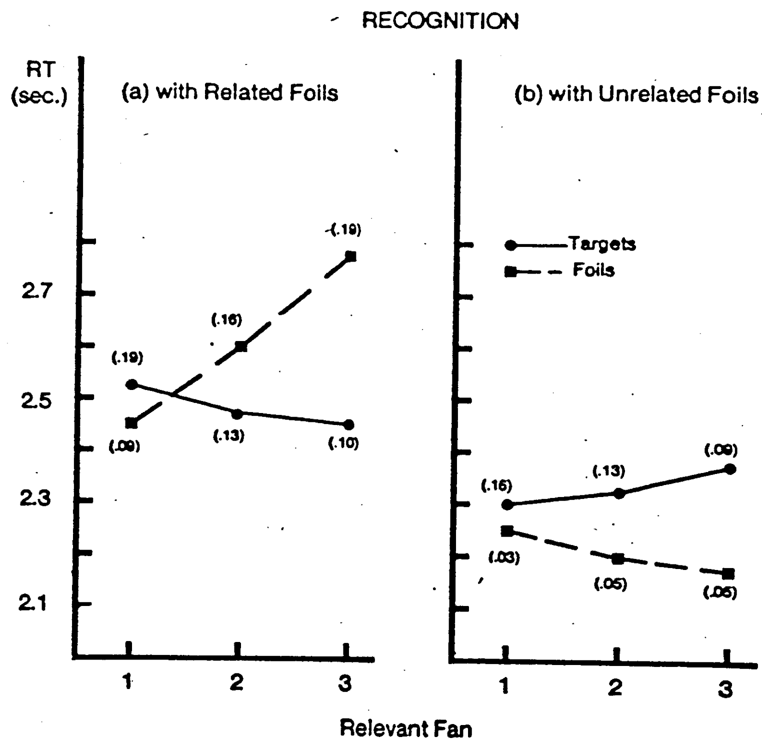

Figure 2. Mean reaction times (RTs) and proportion of errors (in parentheses) as a function of relevant fan in the recognition blocks.

errors and RTs for the related recognition block suggests that the consistency strategy was used some of the time. Use of the consistency strategy would tend to cause errors for foils that were highly thematically related; a "consistent" decision would not be an error for targets though. This explains the opposite trend in errors for targets and foils as a function of relevant fan: The greater the relevant fan, the more consistent the statement. The more consistent the statement, the more likely subjects are to give yes responses, meaning more errors for foils and higher accuracy for targets.

This tendency also explains why there was no effect of relevant fan on RT for targets, $F(2, 64) = 1.14$, whereas there was one for foils. Consistency judgments for foils would be errors and not included in the RTs; RTs for targets would include them and therefore reflect an averaging of facilitation and interference due to fan. Thus, although we did not obtain the usual fan effect on target RT, we are arguing that the pattern of results are consistent with that finding.

Further evidence for this claim is that these results are similar to those of the first experiment of Reder and Anderson (1980). They found no significant fan effect in the test blocks where foils were not thematically related but did for foils in the condition that required discrimination. There, too, accuracy showed the opposite trends for targets and foils in the thematic foil recognition block. Reder and Anderson performed a second experiment to try to force subjects to be more careful in the recognition condition that used thematic foils. By monitoring accuracy, the hope was that subjects would use consistency judgments less in that block. In that way, target RTs would reflect primarily direct retrieval processes and therefore show a fan effect. Using the accuracy-monitoring technique, RTs to studied facts increased with relevant fan. Given the similarity in results between our recognition data (for both related and unrelated foils) and their first experiment, it seems reasonable to assume that if we used the accuracy-monitoring technique, we would obtain results similar to their second experiment.

Consistency judgment task. The pattern of results for relevant fan in the consistency block can be considered a test of our interpretation of the recognition results. Each task reflects a mixture of the direct retrieval and the plausibility judgment strategies. Recognition judgments tested in the presence of unrelated foils are often made using the plausibility judgment strategy; however, we expect that the proportion of trials using that strategy will be even greater when subjects are actually asked to make consistency judgments.

Figure 3(a) plots time to judge whether a probe is consistent with a studied fact as a function of relevant fan. No effect is expected for the foils because a foil to a studied character with three facts on a theme is no more thematically related to studied facts than is a foil to an individual with one fact learned about a theme. That is, if a subject learned about Marty christening ships and the foil was *Marty cooked spaghetti,* the number of facts about Marty christening the ship probably would not affect the time to say *not consistent.* Indeed, we found no effect of relevant fan here, and it is included principally for comparison of overall RT (refer to Footnote 5).

More generally, the pattern in the consistency block is what one would expect given our interpretation of the results in the recognition blocks. There is a significant negative fan effect for thematic statements, $F(2, 64) = 18.03$, $p < .001$, such that subjects are faster to make a consistency judgment, the more facts they know on the relevant topic. The RTs also decrease for studied facts in the consistency block. This effect is smaller than for thematic statements and is not significant, $F(1, 64) = 1.67$. The pattern obtained for accuracy is also what one would expect given our notions and is significant for both the-

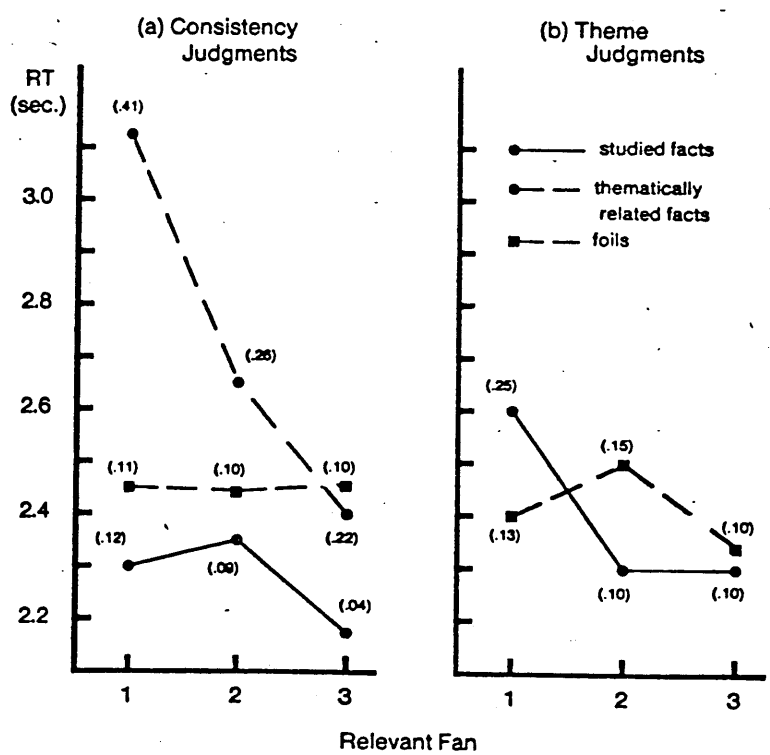

Figure 3. Mean reaction times (RTs) and proportion of errors (in parentheses) as a function of relevant fan (a) in the consistency block and (b) for the theme judgments.

matic statements and studied facts, $F(2, 64) = 14.33$, $p < .001$ and $F(2, 64) = 3.98$, $p < .025$, respectively, such that subjects are more accurate, the greater the relevant fan. The reason that the decrease in RT is smaller and not significant for studied facts, unlike thematic statements, is due to a mixture of strategies. This claim is discussed further in the summary below.

Because the large differences in error rates and in RTs across conditions are in the same direction, there is not a problem of speed–accuracy trade-off (see Pachella, 1974). Nonetheless, it is notable that in the consistency task, the one-fan condition for consistent facts not studied has an error rate of 41%. Indeed, in all of the tasks where consistency judgments can be used at least some of the time, error rates are higher in the low relevant fan conditions. One explanation is that subjects are less likely to form a theme subnode with only one fact. The data of Reder and Anderson also suggested this tendency. When no subnode is formed so that the fact is linked directly to the person node, activation will not intersect at a subnode. Lack of a subnode is less likely to cause an error for presented statements because these probes would be correctly responded to when the direct retrieval strategy is used instead of the official task.

Theme judgments. The theme judgment results were expected to be similar to the consistency judgment results, and they were. Figure 3(b) plots RT to verify that a character was associated with a particular theme, plotted as a function of relevant fan. Here too, the appropriate procedure involves verifying whether a specific subnode is attached to the individual. Again, no effect of relevant fan is expected for foils because the theme mentioned in the foil predicate is unrelated to the themes studied with the probed character (see discussion above and Footnote 5).

The correct character–theme pairings, however, were expected to show a negative fan effect, and they do show a reliable effect, $F(2, 64) = 3.78$, $p < .05$. The error rate also decreases with fan, $F(2, 64) = 10.18$, $p < .01$.

Summary of relevant fan effects. The pattern of data across items and tasks cannot be accounted for by assuming that subjects use one strategy exclusively. Thematic statements give a different RT function and error rate pattern depending on whether they are supposed to be rejected (as foils in the difficult recognition task) or accepted (as consistent statements in the consistency block). Both direct retrieval and consistency strategies are used in this experiment. Given the pattern of error rates, it is extremely unlikely that subjects always use the strategy officially required by the task.

Those data also argue that a different mixture or ratio of the two strategies, direct retrieval and plausibility, are used depending on task characteristics. For example, if the same proportion of trials used the consistency strategy across tasks, then the mean RT to verify studied facts would be the same across tasks, and it is not. Studied facts were judged faster in the block with unrelated foils (2.33 sec) than these same statements tested in the block with related foils (2.62 sec) and were accepted fastest in the consistency block (2.27 sec). This suggests that the more difficult (slower) process of direct retrieval was used most often in the recognition task where foils could not be correctly judged with a consistency strategy; the faster consistency process was used most often when it was the official task.

The notion of a mixture of two strategies being applied to the studied facts easily accounts for the relatively flat functions across relevant fan for those items. Unlike the foils or thematic statements, use of the unofficial strategy on studied facts would not usually cause an error to be made. Instead, the two RT functions, one for direct retrieval process and one for consistency judgment process, would be averaged together. The obtained RT function for thematic foils in the (difficult) recognition task is relatively "pure" as an estimate of the direct retrieval process because use of the inappropriate strategy would cause an error and not be included in RT measure. Similarly, the RT functions obtained for thematic statements in the consistency block are "pure" estimates of the consistency process. Assuming that the studied facts were sometimes answered using one process and sometimes using the other, one can see that a mixture of the corresponding RT functions would produce relatively flat functions. The function goes down for studied facts in the consistency block because the consistency judgment process is used more of the time in that task than in the recognition task.

In addition to predicting when positive, negative, or flat functions will be obtained in the three tasks and predicting the error rates, the model also accounts for the obtained pattern for theme judgments. The negative fan effect is large for correct pairs for the same reason that thematic statements show a large negative fan effect in the consistency task. The judgments cannot be a mixture of recognition and consistency processes; that is, subjects must answer these questions using consistency judgments. Subjects are more accurate for theme pairs than for thematic statements in the consistency task because they are never tempted to try the direct retrieval strategy in the theme judgment task.

One aspect of the results might seem somewhat inconsistent with the proposed model. Given the notion that the consistency judgment process is faster than is the direct retrieval process, one might wonder why thematic statements are judged slower than are studied facts (in the consistency block) since the former items use consistency judgments exclusively. The reason is that, on some portion of these trials, direct retrieval is tried first, and that strategy will never be successful. When direct retrieval is tried first, subjects either respond *no* and make an error, or they then can try the official task. Having tried an unfruitful process first, RTs will be slowed considerably. This however, raises another question: Why do subjects bother to use direct retrieval at all in the consistency task? The answer is not obvious to us, although this same phenomenon was observed in a somewhat different task (Reder, 1982). Our best guess is that at very short delays (between study and test), a direct match of the probe to the studied fact is a fast and easy question-answering strategy and that subjects are reluctant to give up what appears to be an attractive verification process.

Theme Fan

Given that relevant fan produces drastically different results with different tasks, we must show that theme-fan remains unaffected by task characteristics. Theme-fan refers to the number of themes studied with the probed character. The theory predicts that the more themes studied about a character, the longer it would take to make a judgment in any of these tasks. With more themes, there are more links that must divide the activation that is sent from the character node. The amount of activation that reaches a node affects response execution time (to be clarified later). The effects of theme-fan should occur regardless of whether the predicate in the test probe is related to those themes.

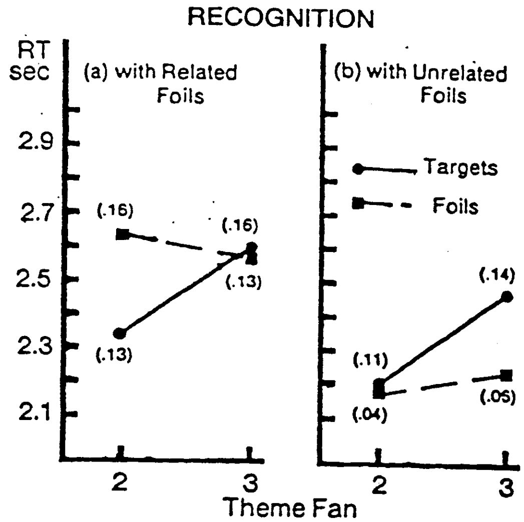

Figure 4. Mean reaction times (RTs) and proportion of errors (in parentheses) as a function of theme-fan in the recognition blocks.

Figure 4 plots the time to make a judgment in the recognition blocks as a function of the number of themes associated with the probed character. Again, the left panel plots the RTs for studied facts and foils tested in the recognition block that used thematically related foils; on the right are RTs for studied facts and foils for the trial block using unrelated foils. Number of themes significantly increases RT for studied facts in both recognition blocks, $F(1, 32) > 5.60$, $p < .05$, for both. Number of themes did not affect RT for the foils in either block, both Fs < 1.0.[6]

[6] We assume this is due to the same strategy hypothesized by McCloskey and Bigler (1980), that is, "the strategy of searching both subsets of facts before responding 'false' represents an attempt to increase accuracy" (p. 262). They found the effect of irrelevant fan greater in experiments that stressed accuracy. Irrelevant fan was constant across levels of theme fan in our experiment.

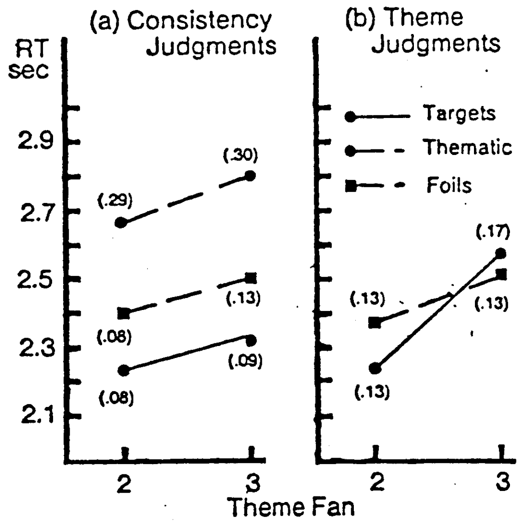

Figure 5. Mean reaction times (RTs) and proportion of errors (in parentheses) as a function of theme fan (a) in the consistency block and (b) for the theme judgments.

Now consider the theme-fan effects for consistency and theme judgments. Figure 5(a) plots RTs to make consistency judgments as a function of number of themes for the three types of statements. Collapsing over type of statement, there is a fan effect for the number of themes, $F(1, 32) = 5.57$, $p < .025$.

Figure 5(b) plots the functions for theme judgments. Because only character name and theme name are presented in a probe, for example, Marty ship christening, there is no distinction between studied facts and thematic statements. Here, too, there is a healthy theme-fan effect such that RTs increase with number of themes, $F(1, 32) = 4.37$, $p < .05$. In all of these comparisons, differences in error rates are small and in the same direction as are the RTs.

Irrelevant Fan

Table 3 indicates the conditions used in an ANOVA to test the effects of irrelevant fan.[7] As in the Reder and Anderson (1980) experiments, RTs were not affected by irrelevant fan, per se. That is, the number of facts about the probed character unrelated to the test probe had no effect on RTs when number of themes was controlled for, all Fs$(4, 128) < 2$ for the four recognition tasks. For consistency judgments, the Fs are less than 1.0. For the theme judgments, there is no effect of irrelevant fan, $F(4, 128) = 2.35$, $p < .10$.[8]

Regression Analyses

Regression analyses were performed to provide a converging measure of which aspects of the experiment affected RT in the various tasks. These analyses were motivated by the nonfactorial nature of the design. (The regression analyses were fit to both individual and average data, but because there was little difference in results, we report the average data.) Four basic models were fit to the each task. Different combinations of parameters were fit by the various models. These parameters included total number of facts studied about a character, relevant fan and irrelevant fan (ignoring irrelevant subnodes), theme fan, and so forth. The model which used only theme-fan and relevant fan as parameters did essentially the best job of accounting for the data, by explaining almost 70% of the RT variance. The fit did not improve appreciably by adding more parameters. Therefore, these two variables appear to be the two most important, consistent with the analyses given previously.

General Discussion

There are three important findings of this experiment. The first, and most important, result is that the effect of relevant fan has been shown to vary in particular ways depending on the task. When subjects must use a process that involves inspecting specific propositions to respond correctly, the more relevant facts to be inspected, the longer the RT. When subjects should use thematic consistency to respond, the usual fan effect reverses, and more relevant facts lead to shorter

[7] We also examined specific contrasts, such as 1-1-0 vs. 1-3-0 and 1-2-1 vs. 1-3-2, within each of the nine tasks and found no evidence for an effect of irrelevant fan. These contrasts are based on a relatively small number of observations, so the overall analyses seemed a fairer test, since we had predicted no differences.

[8] Though it might appear that there is some effect of irrelevant fan here, the means showed no obvious pattern and were not increasing with irrelevant fan. The means for irrelevant fans of 1 to 5 were 2.338, 2.473, 2.317, 2.156, and 2.588 sec.

RTs. The pattern of results was consistent with the view that subjects use both the consistency judgment strategy and the direct retrieval strategy regardless of the official task; however, depending on the task characteristics, the preference (proportion) of use of one or the other strategy changes (as found by Reder, 1982). Evidence that both strategies are used is that error rates for thematic statements are not at chance; these statements must be responded to with different strategies in the different blocks, or performance would be quite poor. Evidence that the ratio of use of the two strategies differs with task is that the mean RT for the presented statements (as well as the pattern of relevant fan effects) varies across tasks.

The second finding is that the more themes studied about an individual, the longer the RT. This result obtains both when subjects may respond using thematic consistency and when they must retrieve a specific proposition. This confirmation and extension of earlier findings by Reder and Anderson (1980) and McCloskey and Bigler (1980) is important in the context of understanding the implications of the first result mentioned. We have shown that the relevant-fact fan effect varies with task and that the theme-fan effect does not.

The third finding is that the number of irrelevant facts has no effect on RT when the number of themes is controlled for.

The theoretical accounts given thus far are extensions and integrations of ideas expressed by Reder and Anderson (1980) and Reder (1982). None of the mechanisms described above can explain why consistency judgments should be facilitated by increased fan. Below we discuss the novel implications of these results for models of memory and propose one means of accounting for these findings that is largely compatible with the theory already presented.

Implications of Findings for Models

Whether or not the reader agrees with the theoretical interpretations given thus far, the results reported in this article impose certain constraints on any model that attempts to account for them. We assume that the learning or study phase of the experiment leads to a single representation of the presented information that is used in all of the later tasks. If this is true, then the very different pattern of results in the different tasks must be explained by changes in the processes that use the information. These processes may be of many different forms such as strategy changes, criterion shifts, or the use of different information on which to base decisions. These process changes must not change the effect of the number of themes, because this effect was the same in all the tasks, but must allow radical changes in the effect of the number of relevant facts. A reasonable way to account for such a pattern is to have the access of the specific propositions be through the themes. When the theme content is sufficient to provide an answer, search for a specific proposition is not necessary.

Specific Proposals

The model proposed in the introduction is one way in which specific propositions can be accessed through themes. That model handles the increase in RT with greater theme-fan, the lack of irrelevant fan effects on RT, and the typical fan effects. The problem with that model is that although it can account for the attenuation of fan effects (as found in recognition tasks that use unrelated foils), it cannot account for the facilitation of RTs with greater fan that occurs in the consistency and theme judgment tasks.

One solution to this difficulty that we considered was to allow the character-to-theme subnode links to vary in strength. The more propositions studied about a particular character on a given theme, the stronger the link. The stronger the link, the shorter the time to traverse it, and hence the shorter the time to activate the information that that character was studied with that theme. This proposal accounts easily for the facilitation effect of relevant fan in the consistency task and still allows the usual fan effect. It assumes that the effect (longer RT) of an additional relevant fact is greater than the effect (time saved) by having a stronger link between the character and theme subnode. This solution accounts for most of the results from our experiment in a way not inconsistent with earlier ACT-based models.

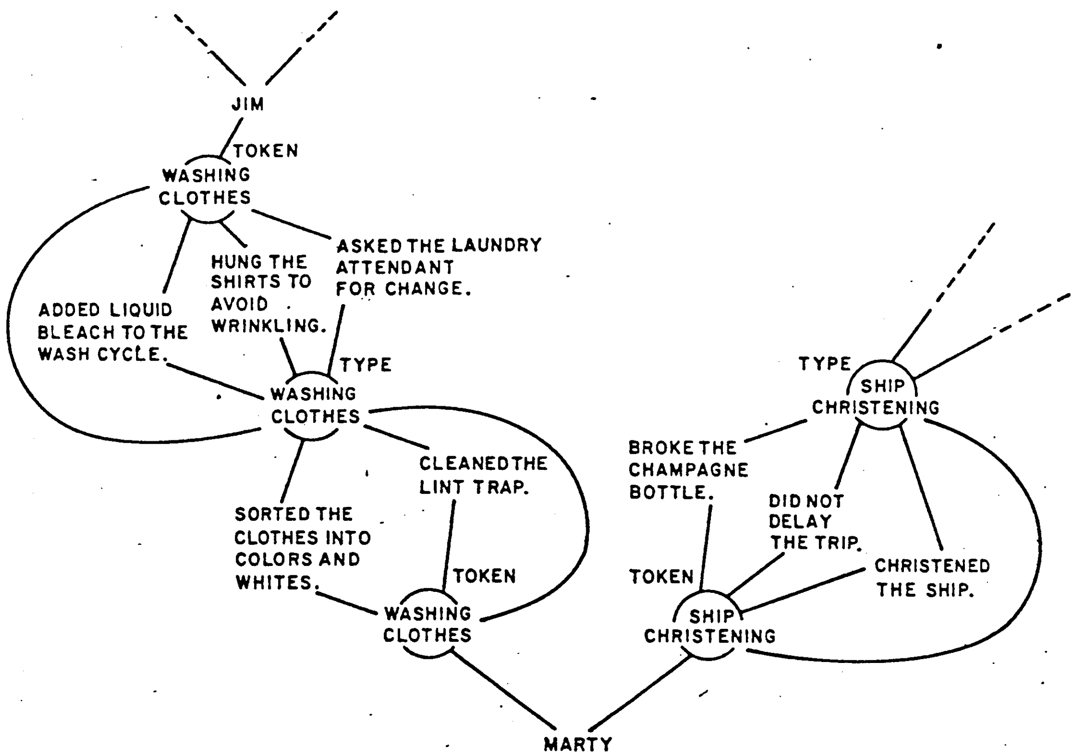

Figure 6. A possible memory representation for related sets of facts about an individual, involving subnode theme *tokens* and theme-*type* nodes.

There is a difficulty with this proposal, however; it predicts an effect of the number of irrelevant facts studied with a character–theme pair. Because the number of facts fanning off a subnode would affect the strength of the character–theme subnode link, the more irrelevant facts studied, the more activation the irrelevant theme would siphon off. With less activation for the relevant character–theme subnode, the RTs would be slowed. Of course, there was no effect of irrelevant fan, such that the number of facts learned about a character on a theme unrelated to the test probe did not matter. In other words, any model attempting to account for theme-fan effects and relevant-fact fan effects that vary with task must also account for no effect of irrelevant fan. Therefore, an adequate model using this framework must have the effect of relevant fan occur between the theme subnode and the facts, not between the character and the theme. So we abandoned the explanation that the number of facts affected the strength of connections between character and subnode.

One memory representation we consider viable is only a slight modification of the memory representation described in Reder and Anderson (1980). It assumes that an individual character's subnode is attached to a generic theme node or *type* node.[9] So, for example, facts about Marty christening a ship would be attached to both the subnode from Marty and to the node that has links to all ship-christening facts, (preexperimental as well as experimental). Figure 6 is a modified version of Figure 1 that partially reflects this type–token distinction. When the test probe is parsed, concepts in the probe related to the theme activate the relevant type node. The type node, like all active nodes, will send out activation to attached nodes, in this case the token nodes (theme subnodes). The amount of activation sent to a particular subnode will depend in part on the number of links connecting the token to the type. In other words, activation reaches the theme subnode from

[9] We thank J. R. Anderson for suggesting this representation.

its type node as well as from the character node. And, importantly, more activation will spread to the subnode from the theme node when the subnode has greater fan. Despite this, the amount of activation of individual propositions is still lower with greater fan because the subnode's activation must be divided between more propositions.

This modified representation by itself does not account for interference or facilitation due to differences in the amount of fan off a node. As before, we must postulate different strategies for different tasks that affect the processes operating on the activated network. As implied earlier, we assume that there are two processes. The ratio or mixture of these two processes varies across tasks and depends on a number of factors, as described earlier (see also Reder, 1982). One type of process is a test for consistency that looks for intersection of activation at a subnode. The person name in the probe activates the corresponding node in memory, and activation spreads out from it. At the same time, the concept words in the predicate activate their corresponding concepts in memory, which also spread activation. The activation from the source nodes in the predicate will intersect at the theme-type node and spread from there to the subnode. Activation will intersect from the person node and the theme-type node if the appropriate theme had been studied with the person.

The second process, the direct retrieval process, tries to match a specific proposition with the test probe. For a proposition in memory to be matched with the test probe, the proposition must be active. The speed of processing for either strategy depends on the level of activation of the elements being tested. A given proposition associated with a person node will be less active when there is greater relevant fan. On the other hand, the activation of the theme node is greater with more relevant fan.[10]

Conclusions

Empirically, the important result is that the effect of relevant fan varies with the task. When subjects must retrieve specific propositions, more relevant facts increases RT, but when subjects may respond by using consistency judgments, more relevant facts decreases RT.

From a theoretical perspective, there are two main points. First, this result appears to resolve the controversy surrounding the effects of fan in thematically related materials. A number of research endeavors using simpler experimental procedures have attempted to account for integrated knowledge effects. Most of these theories are cast into doubt by the results reported here. A representation of memory for integrated material must allow for multiple processes to operate on it that will yield radically different results with different tasks.

The second point to note is that by using methodology that requires subjects to use different processes on the same knowledge structure, we were better able to constrain the class of possible models. The usefulness of complex, within-subjects experiments becomes clearer as more complex theories are developed. As these models attempt to account for a variety of phenomena, the need for this type of experiment that demands many strategies from one set of subjects becomes important.

---

[10] One difficulty with a model such as the one described here is that it is difficult to see whether or not these ideas are sufficient to account for the data. One type of sufficiency test is to implement these ideas in a running computer program. A set of simple productions, using an ACT-based architecture, was developed by Peter Pirolli (Note 1), a student of J. Anderson, to account for this pattern of data. Pirolli's consistency production took less time with greater fan, whereas his direct retrieval production took more time with greater fan.

Reference Note

1. Pirolli, P. Unpublished computer program. Carnegie-Mellon University, January 1982.

References

Anderson, J. R. Retrieval of propositional information from long-term memory. *Cognitive Psychology*, 1974, *5*, 451–474.

Anderson, J. R. *Language, memory, and thought*. Hillsdale, N.J.: Erlbaum, 1976.

Anderson, J. R. *The architecture of cognition*. Cambridge, Mass.: Harvard University Press, in press.

Anderson, J. R., & Bower, G. H. *Human associative memory*. Washington, D.C.: Winston & Sons, 1973.

Clark, H. H. The language-as-fixed-effect-fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, 1973, *12*, 335–359.

Hayes-Roth, B. Evolution of cognitive structures and processes. *Psychological Review*, 1977, *84*, 260–278.

King, D. R. W., & Anderson, J. R. Long-term memory search: An intersecting activation process. *Journal of*

Verbal Learning and Verbal Behavior, 1976, *15*, 587–606.

Lewis, C. H., & Anderson, J. R. Interference with real world knowledge. *Cognitive Psychology*, 1976, *7*, 311–335.

McCloskey, M., & Bigler, K. Focused memory search in fact retrieval. *Memory & Cognition*, 1980, *8*, 253–264.

Moeser, S. D. Recognition processes in episodic memory. *Canadian Journal of Psychology*, 1977, *31*, 41–70.

Moeser, S. D. The role of experimental design in investigations of the fan effect. *Journal of Experimental Psychology: Human Learning and Memory*, 1979, *5*, 125–134.

Pachella, R. G. An interpretation of reaction time in information processing research. In B. Kantowitz (Ed.), *Human information processing: Tutorials in performance and cognition*. Hillsdale, N.J.: Erlbaum, 1974.

Ratcliff, R., & McKoon, G. Does activation really spread? *Psychological Review*, 1981, *88*, 454–462.

Reder, L. M. Plausibility judgments versus fact retrieval: Efficient strategies for question-answering. *Psychological Review*, 1982, *89*, 250–280.

Reder, L. M., & Anderson, J. R. A partial resolution of the paradox of interference: The role of integrating knowledge. *Cognitive Psychology*, 1980, *12*, 447–472.

Schustack, M. W. *Word meaning comprehension: Syntactic and associative effects of sentential context*. Unpublished doctoral dissertation. Carnegie-Mellon University, 1981.

Smith, E. E., Adams, N., & Schorr, D. Fact retrieval and the paradox of interference. *Cognitive Psychology*, 1978, *10*, 438–464.

Thorndyke, P. W., & Bower, G. H. Storage and retrieval processes in sentence memory. *Cognitive Psychology*, 1974, *5*, 515–543.

Appendix A

Instructions

In this experiment you will be learning a large number of facts about a dozen fictitious people. In the first phase of the experiment, one of these people will be displayed on the screen before you, along with all the facts that you are to associate with that person. The amount of time these facts stay on the screen is related to the number of facts you must learn with that person. Then the screen will be erased, and another fictitious individual will be presented along with his or her facts.

It is important to try to learn these facts about these people as quickly as possible. In later phases of the experiment, you will be tested on these facts. How long the experiment takes depends on how quickly you learn the facts.

One suggestion that will probably help you learn the facts quicker is to imagine or visualize a person for each of the people given. Try to create a little story consistent with the facts you learn about each person or try to justify why the specific set of facts is true of the particular person associated with the facts. More than one person will be associated with a given fact. However, it is not a good idea to build connections among the people since no two people have an identical set of facts associated with them.

More instructions will be given to you at various stages throughout the experiment. If you have any questions, or something is bothering you, please go and ask the experimenter. If possible, try not to leave your terminal in the middle of a phase. There are rest breaks between each phase.

Instructions Concerning Topic Judgments

There will be two parts in the reaction time phase that you are now beginning. (The reaction time phase is the last phase and the most critical phase.) The two parts differ with respect to the type of judgments that you will be asked to make. In both parts, a sentence will be displayed on the screen, and you will have to make a judgment about it. The other instruction sheet will explain more about this first type of judgment. Occasionally during this part and the next part, the screen will display the name of one of the professions and a topic instead of a sentence. When this happens, you are to make a judgment as to whether or not that topic has been associated with that person. If you studied that topic with that person, type "d" for yes; if that topic was not studied with that person, type "k" for no.

Please keep you index fingers on d and k at all times so that you can respond as quickly as possible. We are concerned that you be both quick and accurate. It is more important that you be accurate than quick; however, your reaction times will be much more meaningful if you fingers are always ready to respond with the correct key.

If you have any questions, please ask the experimenter.

Instructions for the Consistency Block

In this part of the experiment, you must try to decide whether you think a sentence displayed on the screen is *consistent* with what you studied during the learning phase. Some of these sentences will be ones that you actually studied. These, of course, should be considered consistent. The other sentences will involve facts that you have studied though not with that particular person. Some of these facts, however, are related to the ones that you have studied; these sentences come from the same topic as the facts you studied with that person. These should be considered consistent. Facts that are not from the same topic as other facts associated with a particular person should not be considered consistent.

If a fact is consistent with facts about the person in question, type a "d"; if a fact is not consistent, type a "k." Please remember to keep one index finger resting on the d button, the other on the k. Try to be as quick as you can while remaining accurate.

There will also be topic judgments interspersed with the consistency judgments.

Instructions for the Recognition Block

In this part of the experiment, you will be asked to decide whether or not the sentence displayed on the screen is one that you studied during the learning phase. If so, type a "d" for yes, and if not, type a "k" for no. The sentence will take the form of sentences you studied; that is, the screen will present the name of one of the 12 people studied and one of the facts that you studied about one of the 12 people. In the "no" case, though, the fact will not have been studied with that person.

There are two kinds of recognition blocks, one more difficult than the other. They differ in the nature of the "no" sentences, that is the sentences you did not study and must reject. Suppose you had learned the following three facts about a character called Melvin:

Melvin skied down the slope.
Melvin waited in the lift line.
Melvin waxed his skis.

Compare the following two "no" facts: "Melvin adjusted his skis" and "Melvin cooked spaghetti." Presumably you would find rejecting the first fact more difficult that the second because it is more similar to those facts you would have studied about Melvin. The "Easy Recognition Block" and the "Hard Recognition Block" differ only in the nature of the sentences to be rejected.

Do not forget that there will also be topic judgments interspersed with the recognition judgments. Also, please remember to keep your index fingers on the d and k buttons so that you are ready to respond.

After every 13 trials, you will be told how many errors you made in that interval. We do this to remind you of how important it is to be accurate. Please respond as quickly as you can, but try to stay accurate too. Both are important.

If you have any questions, please see the experimenter.

Appendix B

Data

	Target		Foil			Target		Foil	
Fan	% Cor	Av time	% Cor	Av time	Fan	% Cor	Av time	% Cor	Av time
Recognition judgment: unrelated foils					Recognition judgment: related foils				
110	.939	2.035	.955	2.087	110	.848	2.349	.909	2.581
400	.939	2.156	.977	2.034	400	.939	2.230	.833	2.593
220	.856	2.195	.955	2.162	220	.826	2.441	.833	2.546
310	.904	2.120	.954	2.294	310	.899	2.302	.813	2.825
130	.818	2.221	.985	2.261	130	.818	2.330	.879	2.601
211	.909	2.397	.879	2.247	211	.909	2.625	.924	2.798
121	.758	2.506	.970	2.264	121	.773	2.713	.879	2.284
420	.909	2.179	.962	2.256	420	.939	2.441	.780	2.603
240	.848	2.391	.970	2.264	240	.909	2.294	.803	2.563
330	.954	2.230	.965	2.030	330	.914	2.402	.788	2.740
321	.899	2.570	.919	2.221	321	.889	2.553	.828	2.735
213	.909	2.272	.970	2.144	213	.848	2.456	.833	2.728
132	.848	2.516	.970	2.379	132	.788	2.663	.970	2.310
222	.848	2.398	.960	2.170	222	.848	2.510	.833	2.464

(*Appendix continued*)

Appendix B (*continued*)

	Target		Foil			Stated		Inconsistent		Consistent	
Fan	% Cor	Av time	% Cor	Av time	Fan	% Cor	Av time	% Cor	Av time	% Cor	Av time
	Theme judgments					Consistency judgment					
110	.788	2.528	.871	2.464	110	.939	2.267	.917	2.493	.697	3.077
400	.985	1.873	.924	2.254	400	.955	2.170	.894	2.301	.758	2.070
220	.955	2.141	.818	2.390	220	.939	2.313	.932	2.395	.758	2.606
310	.917	2.020	.909	2.196	310	.949	2.071	.929	2.440	.778	2.275
130	.712	2.536	.856	2.349	130	.864	2.236	.879	2.385	.500	3.033
211	.833	2.631	.848	2.498	211	.924	2.395	.833	2.705	.621	2.869
121	.780	2.402	.826	2.472	121	.833	2.403	.841	2.518	.636	2.969
420	.924	2.357	.894	2.391	420	.947	2.199	.898	2.591	.773	2.235
240	.924	2.275	.848	2.682	240	.864	2.397	.924	2.363	.788	2.729
330	.924	1.909	.947	2.142	330	.960	2.114	.924	2.364	.768	2.408
321	.879	2.661	.864	2.690	321	.960	2.272	.853	2.572	.788	2.511
213	.864	2.469	.848	2.637	213	.894	2.320	.886	2.366	.758	2.652
132	.712	2.976	.924	2.289	132	.879	2.329	.909	2.425	.545	3.539
222	.879	2.070	.899	2.264	222	.914	2.309	.914	2.377	.768	2.490

Note. These values reflect truncations of reaction times above 5 sec. They may not exactly agree with values in particular contrasts described in the article because they do not involve the correction procedure described in Footnote 4. % Cor = percent correct; Av time = average time (in sec).

Received July 21, 1981
Revision received June 21, 1982 ■