Journal of Experimental Psychology: General
1999, Vol. 128, No. 2, 186-197

0096-3445/99/$3.00

The Fan Effect: New Results and New Theories

John R. Anderson and Lynne M. Reder
Carnegie Mellon University

The fan effect (J. R. Anderson, 1974) has been attributed to interference among competing associations to a concept. Recently, it has been suggested that such effects might be due to multiple mental models (G. A. Radvansky, D. H. Spieler, & R. T. Zacks, 1993) or suppression of concepts (M. C. Anderson & B. A. Spellman, 1995; A. R. A. Conway & R. W. Engle, 1994). It was found that the Adaptive Control of Thought—Rational (ACT–R) theory, which embodies associative interference, is consistent with the results of G. A. Radvansky et al. and that there is no evidence for concept suppression in a new fan experiment. The ACT–R model provides good quantitative fits to the results, as shown in a variety of experiments. The 3 key concepts in these fits are (a) the associative strength between 2 concepts reflects the degree to which one concept predicts the other; (b) foils are rejected by retrieving mismatching facts; and (c) participants can adjust the relative weights they give to various cues in retrieval.

The *fan effect* (J. R. Anderson, 1974) refers to the phenomenon that, as participants study more facts about a particular concept, their time to retrieve a particular fact about that concept increases. Fan effects have been found in the retrieval of real-world knowledge (Lewis & Anderson, 1976; Peterson & Potts, 1982), face recognition (J. R. Anderson & Paulson, 1978), the retrieval of schemas (Reder & Ross, 1983; Reder & Wible, 1984), and the retrieval of alpha-arithmetic facts (Zbrodoff, 1995). These effects have been used to study the effects of aging (Radvansky, Zacks, & Hasher, 1996), effects of working memory capacity (Cantor & Engle, 1993), and effects of frontal lobe damage (Kimberg, 1994). These results have played an important role in the original development of human associative memory theory (J. R. Anderson & Bower, 1973) and in the formulation of the Adaptive Control of Thought (ACT) theory (J. R. Anderson, 1976, 1983). The fan effect is generally conceived of as having strong implications for how retrieval processes interact with memory representations. It has been used to study the representation of semantic information (e.g., Myers, O'Brien, Balota, & Toyofuku, 1984; Reder & Anderson, 1980; Reder & Ross, 1983; Reder & Wible, 1984; Smith, Adams, & Schorr, 1978) and of prior knowledge (J. R. Anderson, 1981; Keenan & Baillet, 1980; Lewis & Anderson, 1976; Peterson & Potts, 1982). The fan effect is also clearly related to associative interference in paired-associates learning and to list-length effects in list memory, as J. R. Anderson and Bower (1973) showed. In all cases, one is increasing the associative fan from cues: concepts in sentence memory, stimuli in paired-associates learning, and the list context in list memory.

The major goal of this article is to describe a theory of retrieval that accommodates both the basic fan results and, importantly, those retrieval phenomena that have been thought to involve processes different from those underlying the fan effect. This theory is based on the general Adaptive Control of Thought—Rational (ACT–R) architecture (J. R. Anderson, 1993; J. R. Anderson & Lebiere, 1998). Although the theory has substantial similarities to previous ACT theories, there are also a few significant differences. We describe first the basic fan result, then the basic ACT–R model and how it accounts for this result, and finally an application of this theory to data by Radvansky, Spieler, and Zacks (1993). We then present a new experiment designed to test whether the fan effect could be produced by the suppression mechanism proposed by M. C. Anderson and Spellman (1995) and by Conway and Engle (1994). We report a number of fits of the ACT–R models to data that depend on three key features of ACT–R: its sensitivity to statistical strength of associative relationships, its ability to adjust the weight of cues, and its mechanism for rapidly deciding it has not seen something before.

The Basic Fan Effect

The fan effect is most often demonstrated in recognition memory. Table 1 shows some of the material of the original demonstration (J. R. Anderson, 1974). Participants studied 26 facts about people in locations. Over this set of materials, either one, two, or three facts were studied about each person and location. The term *fan* refers to the number of facts associated with a particular concept. After committing this material to memory, participants were tested in a paradigm in which they had to recognize sentences they had studied (targets) and to reject foil sentences that were novel combinations of the same people and locations (foils). Both the target and the foil probes could be classified according to the number of facts associated with the person and the location

John R. Anderson and Lynne M. Reder, Department of Psychology, Carnegie Mellon University.

This research was supported by Grant SBR-9421332 from the National Science Foundation. We thank Chris Schunn for comments on this article. The models discussed in this article can be found by following the "published models" link from the ACT–R home page: http://act.psy.cmu.edu.

Correspondence concerning this article should be addressed to John R. Anderson, Department of Psychology, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213. Electronic mail may be sent to ja+@cmu.edu.

Table 1
Examples of Experimental Material in the Fan Experiment of J.R. Anderson (1974)

Material studied	Target probes	Foil probes
A hippie is in the park.	3-3. A hippie is in the park.	3-1. A hippie is in the cave.
A hippie is in the church.	1-1. A lawyer is in the cave.	1-3. A lawyer is in the park.
A hippie is in the bank.	1-2. A debutante is in the bank.	1-1. A debutante is in the cave.
A captain is in the park.	—	2-2. A captain is in the bank.
A captain is in the church.	—	—
A debutante is in the bank.	—	—
A fireman is in the park.		—
A lawyer is in the cave.		
—		
—		
—		

Note. Dashes indicate more items.

(these numbers precede the probes in Table 1). Table 2 shows the data from J. R. Anderson (1974) showing how latency increased as the concepts were associated with more facts.

There are a number of noteworthy results in these data: First, despite the inevitable blemishes of noise in the data, there appeared to be an approximately equal increase in latency as the number of facts associated with person or location increased. Averaging over targets and foils, the size of the effects for person were 1.19 s for 1 fan, 1.28 for 2 fan, and 1.30 for 3 fan. For location, there were effects of 1.20 s for 1 fan, 1.25 for 2 fan, and 1.32 for 3 fan. Thus, the effects in this experiment show a rise of a little more than 100 ms on both dimensions. Other experiments have shown that the fan effects can be of different sizes for different types of concepts. In particular, Radvansky et al. (1993) also studied sentences of the form "the object is in the location." When the object was inanimate, they found larger effects of object fan than location fan; when the objects were animate (people), this effect tended to reverse.

The data in Table 2 also illustrate the *min effect,* which has been replicated numerous times (see J. R. Anderson, 1976, for a review). This is the result that latency is more a function of the minimum fan associated with a probe. For instance, participants tend to respond more slowly to the 2-2 fan items than to the 1-3 or 3-1 items even though these have the same total fan. This is used as evidence for some sort of parallel access, with search being more determined by the lower fan concept.

Table 2
Observed Times (in Seconds) to Accept Targets and Reject Foils Data From J. R. Anderson (1974)

Facts about location	Facts about person 1	2	3
Targets			
1	1.11	1.17	1.22
	(1.08)	(1.14)	(1.18)
2	1.17	1.20	1.22
	(1.14)	(1.22)	(1.27)
3	1.15	1.23	1.36
	(1.18)	(1.27)	(1.33)
Foils			
1	1.20	1.22	1.26
	(1.22)	(1.27)	(1.31)
2	1.25	1.36	1.29
	(1.27)	(1.32)	(1.36)
3	1.26	1.47	1.47
	(1.31)	(1.36)	(1.39)

Note. Predictions derived from ACT–R are in parentheses.

Another effect in the data is that participants showed approximately equal fan effects for targets and foils (the target means were 1.16 for 1 fan, 1.20 for 2 fan, 1.26 for 3 fan; the foil means were 1.23 for 1 fan, 1.33 for 2 fan, and 1.37 for 3 fan). Although these data show a somewhat larger fan effect for foils, J. R. Anderson (1976) reviewed data sets in which this was reversed. Almost never is the fan effect for foils twice as large as the fan effect for targets. This is important because it tends to rule out serial, self-terminating search. Generally, it has been difficult to come up with a convincing model for the rejection of foils. The "logical" thing to do might seem to be to do an exhaustive search of the facts that one knows about a concept. Unfortunately, this leads to incorrect predictions of much longer mean times and larger fan effects for foils.

In the ACT* (J. R. Anderson, 1983) theory, it was proposed that participants rejected foils after they had waited more time than it would have taken to accept a target of that fan. This required that participants assess the fan and adjust their waiting times accordingly, an assumption that always strained plausibility. However, it led to the correct prediction that targets and foils would show equal fan effects. It also gave the theory an extra parameter to estimate the mean difference between targets and foils as the extra waiting time.

The following are some of the basic phenomena that a satisfactory theory of the fan effect will need to accommodate:

1. There are effects of the fan of all concepts. However, the relative size of the fan effects can vary as a function of the exact materials.
2. Latency is more a function of the minimum fan associated with the concepts.
3. Fan effects are approximately equal for targets and foils, with mean latencies being only a little slower for foils.

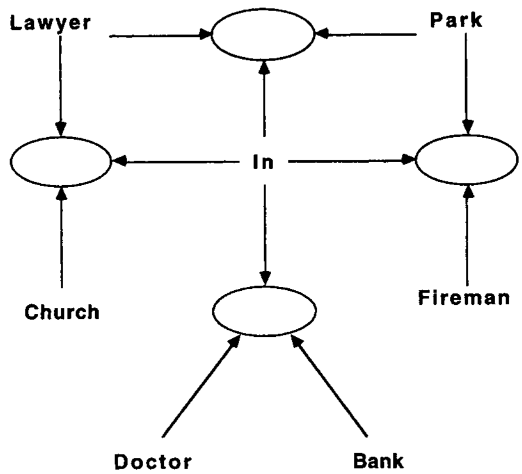

Figure 1. Network representations for four sentences used in the experiment of J. R. Anderson (1974). The sentences are "The doctor is in the bank," "The fireman is in the park," "The lawyer is in the church," and "The lawyer is in the park." Ovals represent facts encoding these sentences and the words represent concepts which are potential sources of activation.

Although these effects seem relatively straightforward, we know of no theory that could provide an adequate account of them. In particular, the ACT* theory and its predecessors could not explain the differential fan effects for different concepts and never really had an adequate theoretical account of foil rejection.

The ACT–R Theory

The basic analysis of the fan effect in the ACT–R theory is similar to the analysis in ACT*. However, the mathematical details are different, reflecting its origins in the rational analysis of J. R. Anderson (1990), and this conveys some advantages to ACT–R. Figure 1 shows a basic network representation of some facts and their associated concepts. As in past ACT theories, activation spreads from presented terms to the connected nodes representing the various facts. According to the ACT–R theory, the latency to retrieve any fact from memory is determined by its level of activation. The basic equation giving the activation, A_i, of a particular fact i is

$$A_i = B_i + \sum_j W_j S_{ji}, \tag{1}$$

where B_i is the base-level activation of fact i and will reflect things such as its past recency and frequency of study. The summation is over the concepts, j, of the probe that provide the sources of activation. In the fan experiment just described, these sources would be the person, location, and preposition *in*. W_j is the amount of attention given to a source j in the probe, and S_{ji} is the strength of association between source j and fact i.

As in past ACT theories, ACT–R's theory of the fan effect turns on the strengths of associations, S_{ji}, between the concepts and the facts. Drawing on the rational analysis of J. R. Anderson (1990), ACT–R has a learning system that produces strengths of association between concepts j and facts i such that

$$S_{ji} = S + \ln [P(i|j)], \tag{2}$$

where S is a constant and $P(i|j)$ is an estimate of the probability of i when j is present. ACT–R bases this estimate on the fan out of j and the proportion of times i occurs when j is present. There are any number of mechanistic learning proposals for how associative strength can reflect probability, including those of Estes (1950) and Rescorla and Wagner (1972). ACT–R is consistent with any of these. Although we complicate this later, as an initial assumption, we set $P(i|j)$ to $1/f_j$ where f_j is the fan associated with the concept j. This reflects the assumption that all facts associated with a concept are equally likely when that concept is present, which is reasonable in most experiments in which all facts are studied and tested with equal frequency. Thus, our strength equation becomes

$$S_{ji} = S - \ln (f_j). \tag{2'}$$

Equation 2′ implies that strength, and hence activation (by means of Equation 1), will decrease as a logarithmic function of the fan associated with the concept. To summarize the mathematics resulting in Equations 2 and 2′, the strengths of associations decrease with fan because the probability of any fact, given the concept, decreases with fan.

We have now shown how activation of nodes depends on associative strengths (Equation 1) and how strengths depend on probability or fan (Equations 2 or 2′). To convert this basic model into a basis for reaction time (RT) predictions requires one additional assumption that is a mapping of activation onto response latency. The basic equation for this in the ACT–R theory is

$$T = I + Fe^{-A_i}, \tag{3}$$

where I is an intercept to represent those activities (such as encoding the probe and generating the response) that do not require retrieval of the critical facts and Fe^{-A_i} is the retrieval time. Thus, retrieval time is an exponential function of activation. The parameter F is a time scale parameter in this equation (i.e., it depends on the units in which time is measured, such as seconds vs. milliseconds).

These equations predict the min effect. To see this, one can substitute Equations 1 and 2′ into Equation 3:

$$\begin{aligned} T &= I + Fe^{-(B_i + \Sigma W_j S_{ji})} = I + Fe^{-[B_i + \Sigma W_j [S - \ln(f_j)]]} \\ &= I + F[e^{-(B_i + \Sigma W_j S)}][e^{\Sigma W_j \ln(f_j)}] \\ &= I + F'\left(\prod_j f_j^{W_j}\right), \end{aligned}$$

where $F' = Fe^{-(B_i + \Sigma W_j S)}$. If one assumes that the weights

W_j are equal to a single W, this becomes

$$T = I + F'\left(\prod_j f_j\right)^W.$$

Note that this makes RT a function of the product of the fans, Πf_j, hence producing the min effect (because the product of a set of numbers with a constant sum is maximum when the numbers are equal, such as $2 * 2 > 3 * 1$ although $2 + 2 = 3 + 1$).

An assumption introduced by J. R. Anderson, Reder, and Lebiere (1996) to account for capacity limitations in retrieval is that

$$\sum W_j = 1. \quad (4)$$

This assumption is used to constrain the weights W_j throughout this article. This prevents the system from having unbounded activation.

Application to the Original Fan Experiment

With the mathematics of the ACT–R theory now specified, we illustrate its application to the experiment in Table 2. Given the representation in Figure 1, we assume that the sources (the *j*s) of activation are the person, the preposition *in*, and the location. Assuming an equal division of activation among these sources, this makes each W_j = one third.[1] Later, in discussing the Radvansky et al. (1993) data, we consider what happens if the W_j varies for person and location but the sum of the W_j retains the constraint of Equation 4.

The time to recognize a fact will be determined by the activation levels of the fact. Assuming that each of the three sources gets a weighting of one third, the activation of the target fact becomes (based on Equation 1)

$$A_{\text{target}} = B + .333(S_p + S_{\text{in}} + S_l),$$

where B is the base-level activation of the target, S_p is the strength of association from the person, S_{in} is the strength from *in*, and S_l is the strength from the location to the target. The preposition *in*, unlike the person and location sources, is used in all facts and therefore does contribute equal strength to all conditions. For simplicity, we can therefore drop this constant value from our activation equation:

$$A_{\text{target}} = B' + .333(S_p + S_l), \quad (5)$$

where $B' = B + .333S_{\text{in}}$. Note that $S_p = S - \ln(f_p)$ and $S_l = S - \ln(f_l)$ from Equation 2′, where f_p is the fan associated with the person and f_l is the fan associated with the location.

J. R. Anderson (1993) proposed that foils are recognized by retrieving some proposition that involves either the person or location. If the retrieved proposition does not match the target, the participant will respond "false." The activation of the nonmatching facts will be either

$$A_{\text{foil}} = B' + .333S_p \quad (5')$$

if the proposition involves the person, or

$$A_{\text{foil}} = B' + .333S_l \quad (5'')$$

if the proposition involves the location.

The mismatching facts have one less source of activation than do matching facts. Thus, the activation of a mismatching fact will be lower than that of a matching fact and the system will retrieve a matching fact when there is one. For simplicity, we assume that half the time participants retrieve a mismatching fact from the person and that half the time they retrieve a mismatching fact from the location. These activation equations can be converted into latency predictions using Equation 3.

Table 2 shows the fit of ACT–R to the original fan experiment. Because we cannot separately estimate F and B', we set B' to zero and just estimated F. Thus, the free parameters in fitting this experiment are the intercept I, the latency scale F, and the initial strength S. These parameters were estimated at I = 845 ms, F = 613 ms, and S = 1.45. The correlation with the data is .87, which compares favorably with the original ACT model (J. R. Anderson, 1976), which had a correlation of .91 but had separate parameters giving different intercepts for targets and foils. An important aspect of the current model is its ability to account for the fan effects for targets and foils and the mean latency difference between targets and foils, all in terms of the same activation parameters.

To illustrate how these predictions were obtained, consider two cases from Table 2.

1. *The 2-2 target:* The strength of association to both the person and location in this case is $1.45 - \ln(2) = .75$. The activation in this case will then be $.333(.75 + .75) = .50$. From Equation 3, we derive the latency as $.845 + .613 * e^{-.50} = 1.216$ s.

2. *The 1-3 foil:* If the participant retrieves from the 1-fan person, the activation is $.333[1.45 - \ln(1)] = .48$. If the retrieval is from the 3-fan location, it is $.333[1.45 - \ln(3)] = .12$. The first activation gives a predicted latency of $.845 + .613 * e^{-.48} = 1.224$ s. The second activation gives a predicted latency of $.845 + .613 * e^{-.12} = 1.389$ s. The average of these two is 1.307 s.

The fit to this model was obtained using the Solver function in the Excel library. This Excel model and a running ACT–R 4.0 (J. R. Anderson & Lebiere, 1998) model of this task can be obtained by visiting the ACT–R Web site at http://act.psy.cmu.edu/ and following the "published models" link.

There are three significant differences between this model and earlier ACT models. One is the integrated treatment of targets and foils. The second is the introduction of W_js to

[1] It might be more plausible to assume *in* receives lesser weighting, but the predictions in this article would not change if we assumed a smaller fraction of the weight for *in* as long as it was constant. Again, it might be more realistic to assume that there is some experimental context as part of the associative structure in Figure 1 that serves as an additional activation source, but the predictions would not change.

weight sources of activation. The third point is that ACT–R uses a past history of use to set the strengths of association to reflect the probability of a fact occurring given the presence of the concept. The second difference becomes important in the next section when considering the data of Radvansky et al. (1993). The third difference becomes important in examining the suppression proposals of M. C. Anderson and Spellman (1995) and of Conway and Engle (1994).

Differential Fan Effects

Radvansky et al. (1993; see also Radvansky & Zacks, 1991) reported a series of experiments in which the fan of either the object or the location was manipulated from one to three and the fan of the other was held constant at one. In their Experiments 1 and 2, in which inanimate objects were placed in locations, they found large effects of object fan and weak effects of location fan. In their Experiments 3 and 4, in which people were placed in locations, these effects were somewhat reversed (i.e., larger effects of location fan and weaker effects of person fan). In their Experiments 5 and 6, in which small locations were used, the effect of location fan was even larger and the effect of person fan was even weaker. Figure 2 shows their results averaged over targets and foils.[2] Radvansky et al. interpreted their results in terms of situation models (Zwaan & Radvansky, 1998). They claimed that participants organized their memory into location-based situation models (where all the objects were in the same location) for Experiments 1 and 2, or that they organized them according to person-based situation models (where the history of a person was represented going from location to location) for Experiments 3 and 4 and particularly for Experiments 5 and 6, where the locations typically held a single person. If all the objects were in one location or all of the locations were associated with a single person, there was only one model to be activated and participants did not have to search through multiple models. They ascribed the fan effect to the need to search through multiple models. As they acknowledged, it is a bit mysterious why there would not be an effect of the number of things in a model (objects in a room or places a person goes to). It should take longer to sort through more things. In fact, there were weak effects of fan for the other dimension. However, they suggested that such size-of-model effects are much less than the effects of having to retrieve multiple room models in Experiments 1 and 2 or to retrieve multiple person models in Experiments 3–6.

The ACT–R model can accommodate these data by varying the weightings (W_j) given to the object and location cues. We fit the same model described earlier to fit J. R. Anderson (1974) in order to fit the Radvansky et al. (1993) data but allowed ourselves to differentially weight the object and location fan, constraining the two to still add to .67 (keeping the weight for the preposition *in* to be .33 in accordance with Equation 4). Figure 2 shows the fit of the ACT–R model to the three data sets. The weighting of the location was .14 in Experiments 1 and 2, which used inanimate objects (and the weighting of the object was .53); the weighting of the location was .44 in Experiments 3 and 4, which used persons (and the weighting of the person was .23); and the weighting of the location was .55 in Experiments 5 and 6, which used small locations (and the weighting of the person was .12). The other parameters were S = 1.28 and F = 1,143 ms[3] for both experiment sets and I = 840 ms for Experiments 1 and 2, I = 798 ms for Experiments 3 and 4, and I = 918 ms for Experiments 5 and 6 (different intercept estimates were obtained to deal with the fact that participants had overall different mean times in the different sets of experiments). In fitting these data, we predicted data separately for targets and foils even though Figure 2 shows the results aggregated over this dimension. Although we estimate different weighting functions for the experiments, the value of F is constant throughout making this a nontrivial test of ACT–R. Basically, we are showing that there is the same amount of associative interference in all these experiments and that it just distributes itself differently between the two concepts.

It is interesting that the ACT–R explanation of the effect is almost the opposite that of Radvansky et al. (1993). They claimed that the concept that shows no fan effect is being used to organize the stimuli into a single model. We question whether it is serving such an organizational role. Rather, we suggest that this concept (for whatever reason) is being largely ignored in the weightings. Instead, we propose that the concept that shows the large fan effect received most of the attention (i.e., had a larger W_j).

What seems indisputable from the Radvansky et al. (1993) experiments is that relative fan effects can vary as a function of material. However, two things are still very much at issue. The first is whether this occurs because participants (a) organize their knowledge at study around one of the concepts into a mental model, as Radvansky et al. argued, thereby avoiding the effects of fan from that concept or (b) focus on one concept at testing, thereby manifesting an effect of that concept's fan. In fact, differential fan effects have been shown for the same materials depending on testing procedures (J. R. Anderson, 1974; King & Anderson, 1976; Reder & Anderson, 1980; Reder & Ross, 1983; Reder & Wible, 1984), which suggests that the effect may be at retrieval, not study. For example, Reder and Anderson found strong fan effects in some trial blocks and no fan effects in others. Importantly, these effects were obtained or not obtained with the same study materials and the same participants. What differed were the foils at testing such that different foils enabled different strategies, as predicted. The important point is that we know that these were not effects due to organization of the information at study because the study material was the same and the participants were the

[2] We would like to thank G. Radvansky for making his data available.

[3] F has a value that is larger here than in the previous experiment. Because the base-level activation, B, is absorbed into the F parameter, the value of F will reflect the amount of practice with the material. The reason why F is different for these experiments in comparison with J. R. Anderson's (1974) is that the earlier study had required more practice with the facts.

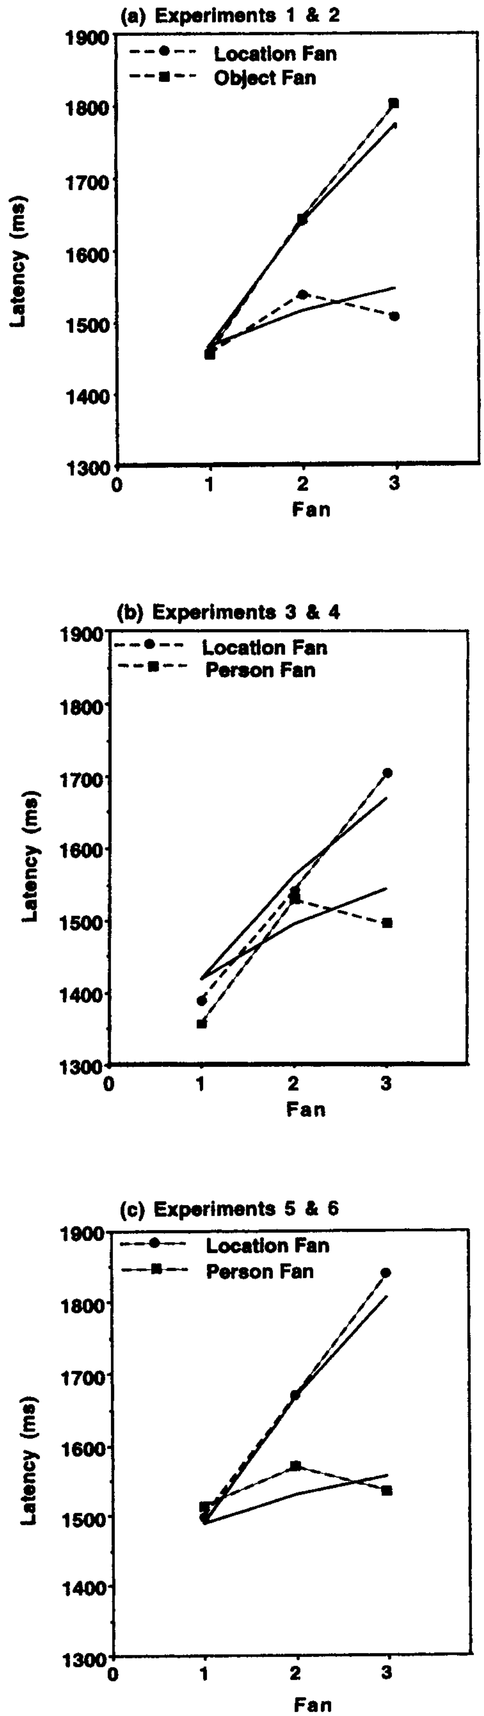

same. These effects were caused at testing due to strategy differences in the evaluation of the study material.

The other unresolved issue is specifying the critical material difference that produced these effects in the Radvansky et al. (1993) experiments. Radvansky et al. offered an explanation in terms of situation models, but the type of situation model differed from experiment to experiment (e.g., whether locations or other concepts provided the models). They provided no converging evidence that this was the way their participants were organizing the materials. Note that there are a number of other differences among their inanimate objects, people, and locations. In a rating study of their material, we found that their inanimate objects had a mean concreteness rating of 6.29, their large locations had a mean rating of 5.49, their small locations had a mean rating of 4.70, and their persons a mean rating of 4.70. Not all of their terms were single words and could be as long as four words ("back room's tanning bed"). Their inanimate objects had a mean length of 1.83 words, their people 1.0 words, their large locations 1.83 words, and their small locations 2.75 words. In terms of our ACT–R model, participants gave more attentional weighting (the W_j parameter) to longer or more concrete phrases, which seems roughly intuitive. Paivio (1971) summarized evidence that concrete words are better cues and multiple-word phrases seem to actually provide multiple cues for retrieval, which, in aggregate, should receive greater weighting.

We do not claim to know what the actual material variable is that controls the relative fan effects. However, we do claim that these can be explained by attentional effects at retrieval and that differences among materials is roughly consistent with such an explanation. We emphasize that the original results of J. R. Anderson (1974) involved symmetrical fan effects for persons and locations and so asymmetrical results are not a necessary feature of this procedure and, certainly, even as Radvansky et al. (1993) showed, sometimes large location fan effects can be obtained.

Fan Effect: Interference or Suppression?

The basic analysis of the fan effect in ACT–R is that as more facts are associated with a concept, the weaker the strengths of associations from the concepts to these facts become and less activation is spread to the facts. More specifically, ACT–R makes the strength of association, S_{ji}, reflect the probability in the past that fact *i* occurred when concept *j* was present (Equation 2). As more facts are known or associated with a concept, the less likely that any one fact will occur in the concept's presence, weakening the strength of association for all facts with that concept. This is an analysis that attributes the interference effects to the associa-

Figure 2 (opposite). Data from Radvansky, Spieler, and Zacks (1993) and predictions of the theory. a: From their Experiments 1 and 2, which had objects in locations. b: From their Experiments 3 and 4, which had persons in locations. c: From their Experiments 5 and 6, which had persons in small locations. Dotted lines connect the data points, whereas solid lines are the predictions of the ACT–R theory.

tive link between the cue and the target fact. Recently, however, M. C. Anderson and colleagues (Anderson, Bjork, & Bjork, 1994; Anderson & Spellman, 1995) and Conway and Engle (1994) argued that there is a suppression or inhibition of memories associated with a particular concept rather than a weakening of specific links. We consider here whether this inhibiting mechanism could be the cause of the fan effect. First, however, we review the paradigm introduced by M. C. Anderson and Spellman and their results. This serves as the basis for an experiment in the fan paradigm testing for suppression.

M. C. Anderson and Spellman's (1995) basic paradigm is one in which participants practice associations to a category. For example, participants might have associated both *blood* and *tomato* with the category *red* but received differential practice, such that *blood* got extra practice. Not surprisingly, this improved the recall of *blood* to *red*. It also lowered the recall of *tomato* to *red*. This inhibitory effect is also predicted by the ACT–R theory. Differential practice strengthens the favored association such that more activation will spread to the more practiced association and less to the less practiced association.[4]

M. C. Anderson and Spellman's (1995) more surprising results concern performance on another related category, *food*, which did not receive differential practice. They did not actually test the recall of *tomato* to that category, but they found that the recall of items related to *tomato*, such as *strawberry* (related because both are red foods), were inhibited compared with the recall of unrelated food items such as *crackers*. They argued that this was because participants actively inhibited or suppressed all red foods. They did so to prevent *tomato* from intruding when they were producing recall of *blood* to *red*. Thus, the fundamental claim of a suppression model is that an entire category of items can be made less available when they share a suppressed element. This suppression result is not predicted by ACT–R because there was no impact of the differential practice of red things on the association between *food* and *strawberry*. That is, it is only the links from *red* to other concepts that should suffer because of the strengthening of the association of *red* from *blood*.[5]

M. C. Anderson and Spellman's (1995) result was obtained in a paradigm different from the one used for the standard fan effect. Their paradigm involves the recall, not the recognition, of facts that were studied in the experiment, organized by real-world categories, and cued for recall by category. We wondered whether similar results could be obtained in a fan experiment. To our knowledge, this has never been examined systematically. It is possible that our attribution of fan effects to decrements in associative strength is actually wrong and that the correct attribution is to a suppression process. That is, perhaps the reason that participants do poorly on high-fan items is because of suppression of facts involving a high-fan concept rather than associative interference. Therefore, for instance, to prevent "The hippie is in the church" from intruding when retrieving "The hippie is in the park," participants could suppress facts associated with *church*. Similarly, retrieval of "The hippie is in the church" should suppress facts associated with *park*. The net effect of such mutual suppressions would be to make all facts associated with *hippie* less available. Although M. C. Anderson and Spellman (1995) did not explicitly address fan effects, Conway and Engle (1994) proposed that fan effects are a result of suppression. Inspired by M. C. Anderson and Spellman's paradigm, we modified the standard fan paradigm to test whether fan effects are due to suppression.

The basic design of our experiment involved participants studying person–location sentences. For half the participants, the fan of the person was held at 2, whereas location fan was either 2 or 4. For the other half of the participants, this was reversed. The structure of the material in the case where location fan was manipulated is illustrated in Table 3. Table 3 represents half of the material of the experiment. In addition to the 24 targets in Table 3 (sixteen 4-fan and eight 2-fan), there was another set of 24 identically structured targets. In all, then, participants studied 48 items.

The third column of Table 3 shows the frequency with which the various facts were tested in a block of the experiment. In a block of the experiment, there were 144 trials involving 72 targets and 72 foils. Most targets and foils were tested once per block. However, there was one type of target fact (designated *A*) that received five tests per block. The foils were designed so that each term (location or profession) was tested as often as a foil as it was as a target. To achieve this, we had to combine some items from the other set of 24. These are the foils labeled *K* in Table 3. For instance, "teacher" in the "teacher–tower," 2-fan pair was studied with targets in the other set of 24.

The critical predictions concern the targets. These predictions are illustrated in Figure 3. All theories predict Target A ("biker–tower" in the figure) to be facilitated because of its repetition. In the ACT–R theory, this is because of the strengthening of the links and we have illustrated this with thickened lines. All theories predict Facts B and C to do worse. In ACT–R, this is because of associative interference from the links encoding A. We have indicated this by dotted lines to denote the weaker associations. In a suppression account of the fan effect, this is because all facts involving the interfering concepts, *factory* and *doctor*, are suppressed to prevent C ("biker–factory") and B ("doctor–tower") from intruding on retrieval of A ("biker–tower"). We have indicated this suppression by shading concepts factory and doctor. Therefore, suppression theory, but not associative interference theory, also predicts worse performance on Fact D because it contains the suppressed *factory* and G because it contains the suppressed *doctor*. Thus, we refer to Target A as defining the facilitation condition, Targets B and C as defining the interference condition, and Targets D and G as defining the suppression condition. The remaining items (E,

---

[4] From an adaptive perspective, systems should favor more practiced associations (e.g., red has become a better predictor of *blood* and a poorer predictor of *tomato*).

[5] Therefore, ACT–R would not even predict a suppression of *tomato* in the context of *food*, but M. C. Anderson and Spellman (1995) did not test that prediction.

Table 3
Structure of Experimental Material: One of Two Sets a Participant Studied

Condition	Pairings	Number of tests (per block)
Targets		
Fan 2		
A (facilitation)	biker–tower	5
B (interference)	doctor–tower	1
C (interference)	biker–factory	1
D (suppression)	writer–factory	1
E (control)	writer–desert	1
F (control)	monk–desert	1
G (suppression)	doctor–bank	1
H (control)	monk–bank	1
Fan 4		
A (facilitation)	cowboy–tunnel, beggar–tunnel	5
B (interference)	soldier–tunnel, singer–tunnel	1
C (interference)	cowboy–cage, beggar–cage	1
D (suppression)	typist–cage, artist–cage	1
E (control)	typist–library, artist–library	1
F (control)	grocer–library, reporter–library	1
G (suppression)	soldier–windmill, singer–windmill	1
H (control)	grocer–windmill, reporter–windmill	1
Foils		
Fan 2		
I (mixed)	writer–tower	1
J (mixed)	monk–tower	1
K (high)	teacher–tower	4
L (low)	doctor–factory	1
M (low)	monk–factory	1
N (mixed)	biker–desert	1
O (low)	doctor–desert	1
P (mixed)	biker–bank	1
Q (mixed)	writer–bank	1
Fan 4		
I (mixed)	typist–tunnel, artist–tunnel	1
J (mixed)	grocer–tunnel, reporter–tunnel	1
K (high)	fireman–tunnel, sailor–tunnel	4
L (low)	soldier–cage, singer–cage	1
M (low)	grocer–cage, reporter–cage	1
N (mixed)	cowboy–library, beggar–library	1
O (low)	soldier–library, singer–library	1
P (mixed)	cowboy–windmill, beggar–windmill	1
Q (low)	typist–windmill, artist–windmill	1

Note. This material is illustrated with the location fan manipulated, but for half the participants the person fan was manipulated.

F, and H; not shown in the figure) define the control condition and should not enjoy facilitation or suffer interference or suppression under any theory. A suppression explanation would not predict a performance difference between interference and suppression facts. In contrast, ACT–R predicts that the interference facts will be worse than the suppression facts and that suppression facts should be equal to control facts.

With respect to foils, there are three meaningful classifications: (a) Foil K involves both concepts that occur in frequently practiced (five times per block) facts; (b) Foils L, M, O, and Q involve concepts that occur in only infrequently practiced facts; and (c) Foils I, J, N, and P involve one concept that occurs in a frequently practiced fact and the other one that does not. These define the high, low, and mixed conditions. ACT–R predicts that participants reject foils by retrieving a fact associated with one of the concepts in the foil and detect a mismatch with the probe. The speed with which such a mismatching fact can be retrieved depends on the strength of the association, which depends on its prior practice. Thus, the expectation of the ACT–R theory is that participants should be fastest to reject in the high condition because there are high-frequency facts that can be retrieved, intermediate in the mixed condition because there is one high-frequency fact, and slowest in the low-frequency condition for which there are no high-frequency facts. This is a unique prediction of the ACT–R model of the fan effect and serves to distinguish it from earlier ACT models. To our knowledge, other theories do not predict any effect of this factor either.

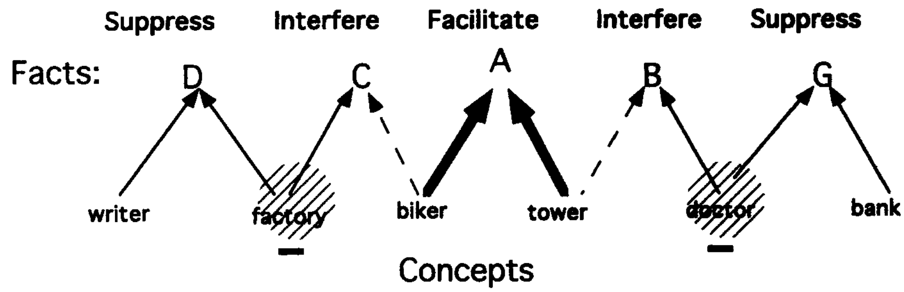

Figure 3. A representation of the critical relationships among the facts and concepts in Table 3. Bold lines represent potentially strong links, dotted lines represent potentially weak links, and shaded concepts are potentially suppressed concepts.

There are numerous procedural details that differ between a fan experiment and M. C. Anderson and Spellman's (1995) paradigm. This includes the use of the episodic rather than semantic material, recognition rather than recall, latency rather than accuracy, simultaneous presentation of material rather than sequential presentation of suppressing information followed by testing of the suppressed information, and so on. Therefore, it is by no means obvious that Anderson and Spellman would predict suppression in this paradigm. Conversely, failure to find suppression in this paradigm does not invalidate their proposal for suppression in their paradigm. The goal of this experiment was to determine whether suppression could be the cause of the fan effect in the typical paradigm where the fan effect has been demonstrated. Also, because the manipulation involved the frequency of various facts, we could test whether participants were sensitive to the statistics of presentation as ACT–R predicts. This test was crucial because such statistical sensitivity is a fundamental part of the ACT–R architecture.

Method

Materials

Each participant had 48 target sentences and 54 foils constructed with respect to the design specified in Table 3. For one group of participants, the people names had their fan manipulated and the locations had a constant fan of 2, whereas for the other group of participants, this was reversed. To construct these targets and foils required a vocabulary of 24 persons and 24 locations. These were chosen to be common nouns, with lengths between four and nine letters with distinct spellings and meanings. Each participant had a different randomly generated set of targets and foils from this vocabulary.

Procedure

The entire experiment was presented on the Macintosh computer with the sentences displayed on the screen and participants responding by keypress. The experiment consisted of three phases. First, participants studied the 48 sentences presented in random order at the rate of 5 s per sentence. The sentences were presented in the form "The person is in the location."[6]

The phase after studying the sentences involved a two-pass dropout procedure. During each pass, participants were presented with all possible questions, both of the form "Where is the person?" and "Who is in the location?" Participants had to type in all locations associated with that person or people in that location, respectively. If they were able to recall all the answers to a question, that question was dropped out of the pass. If they could not, they were shown the correct answer and the question was repeated after all other questions in that pass had been asked. This continued until all questions had been answered correctly one time. A second pass was then administered in the same fashion. Participants averaged 2.5 tests per question to meet the criterion of one correct answer on the first pass and 1.3 questions to meet the criterion in the second pass.

The third phase of the experiment was the critical RT phase in which the targets and foils were repeatedly presented according to the frequencies shown in Table 3. Altogether, there were 144 questions, retested in three successive blocks, for a total of 432 trials. Within a block, the order of the probes was randomized. Each block was broken into four miniblocks of 36 trials, with an optional break after each miniblock.

Each RT trial involved the following sequence: A sentence was presented of the form "The person is in the location." The participant would either press a *K* key for yes or a *D* key for no to indicate whether the sentence had been studied. Feedback in the form of "correct" or "incorrect" was then presented for 1 s. The next trial began after this 1-s feedback. At the beginning of each miniblock of 36 trials, the participant hit a key to begin.

Participants were given points to motivate them to perform their best in the RT phase. Correct responses got an automatic 2 points. If participants were correct, they also got an extra 1 point for each 100 ms that they were faster than 1,500 ms. Incorrect responses cost 20 points and there was no speed reward. These points were converted at a rate of half a penny per point. This was combined with a base pay of $5, for an average pay of about $12 for an experiment that lasted approximately 1.75 hr.

Participants

Sixty-one participants were run from the local Carnegie Mellon and high school populations (16 years of age and older). Two participants were eliminated because of abnormally long RTs (more than twice the mean of the other participants). Of the remaining participants, there were 31 participants in the condition

[6] Note that the person and location were preceded by a definite article. Radvansky, Spieler, and Zacks (1993) did not find an effect of choice of definite or indefinite article.

in which person fan was manipulated and 28 in the condition in which the location fan was manipulated.

Results and Discussion

Separate analyses of variance were performed on targets and foils for latency and accuracy. The factors in these analyses were fan (2 or 4), group (one group had person fan manipulated and the other group had location fan manipulated), type of probe (for target: facilitation, interference, suppression, and control; for foil: high, mixed, and low), and block of the experiment (three values).

The overall results for targets, averaged over block and group, are presented in Table 4. The significant main effects were as follows:

1. Fan for both latency, $F(1, 57) = 7.28$, $p < .01$, $MSE = 248{,}579$, Fan 2 = 1,441 ms, Fan 4 = 1,513 ms; and error rate, $F(1, 57) = 13.71$, $p < .001$, $MSE = 0.0146$, Fan 2 = 6.8% errors, Fan 4 = 9.2% errors.

2. Probe type for both latency, $F(3, 171) = 13.54$, $p < .001$, $MSE = 180{,}379$, facilitation = 1,376 ms, interference = 1,580 ms, suppression = 1,471 ms, and control = 1,479 ms; and errors, $F(3, 171) = 17.67$, $p < .001$, $MSE = 0.0146$, facilitation = 6.1% errors, interference = 12.3% errors, suppression = 6.5% errors, and control = 7.1% errors. The interaction between fan and probe type was not significant for latency, $F(3, 171) = 1.39$, $p > .25$, $MSE = 131{,}098$, or for error rate, $F(3, 171) = 1.69$, $p > .15$, $MSE = 0.0166$.

3. Block for latency, $F(2, 114) = 10.86$, $p < .001$, $MSE = 3{,}313{,}355$, Block 1 = 1,576 ms, Block 2 = 1,444 ms, and Block 3 = 1,410 ms; but not errors, $F(2, 114) = 0.35$, $MSE = 0.0091$.

There were no significant effects of whether person fan or location fan was manipulated. There were also no significant interactions in these data. Specific t tests were performed to test the predictions of the theories with respect to latency for the different probe types. Interference was significantly worse than suppression, $t(171) = 3.52$, $p < .001$, there was no difference between suppression and control, $t(171) = 0.26$, and control was worse than facilitation, $t(171) = 3.32$, $p < .001$. Note that there was no evidence in this experiment of a suppression effect.

The overall results for foils are also presented in Table 4. The significant main effects were as follows:

1. Fan for both latency, $F(1, 57) = 9.51$, $p < .01$, $MSE = 279{,}935$, Fan 2 = 1,559 ms, Fan 4 = 1,658 ms; and error rate, $F(1, 57) = 24.68$, $p < .001$, $MSE = 0.0138$, Fan 2 = 9.2% errors, Fan 4 = 12.7% errors.

2. Probe type for both latency, $F(2, 114) = 4.44$, $p < .05$, $MSE = 161{,}777$; high = 1,558 ms, mixed = 1,625 ms, and low = 1,643 ms; and errors, $F(2, 114) = 11.00$, $p < .001$, $MSE = 0.0152$; high = 8.7% errors, mixed = 11.0% errors, and low = 13.1% errors.

3. Block for both latency, $F(2, 114) = 12.98$, $p < .001$, $MSE = 301{,}865$, Block 1 = 1,724 ms, Block 2 = 1,585 ms, and Block 3 = 1,516 ms; and errors, $F(2, 114) = 4.14$, $p < .05$, $MSE = 0.0163$; Block 1 = 9.5% errors, Block 2 = 11.2% errors, and Block 3 = 12.2% errors. Note that the error effect for blocks was weak and was in the opposite direction of the latency effect.

The foils also did not show a significant effect of the type of item whose fan was manipulated, and the data in Table 4 are collapsed over this variable. There was a Block × Probe Type interaction that was significant only for errors, $F(4, 228) = 3.84$, $p < .01$, $MSE = 0.0078$. The differences among the probe types increased as the number of trials increased. That is, the effect of differential frequency of the targets on the foils increased with practice. No other interactions were significant for either dependent measure.

The latencies and errors generally correlated (the overall correlation between latency and errors in Table 4 is .88). We focus mainly on the latency data. There are three significant conclusions with respect to the theoretical issues that motivated this experiment: First, there was no significant effect of item whose fan was manipulated (the group variable in our design). For our targets, the mean fan effect (Fan 4 − Fan 2) for locations was 80 ms, whereas it was 73 ms for people, $F(1, 57) = 0.01$, $MSE = 248{,}579$. For foils, the mean fan effect was 133 ms for locations, whereas it was 67 ms for people, $F(1, 57) = 1.03$, $MSE = 279{,}935$. Although the effects were not significant, they were in the direction (greater fan effect for locations) as reported by Radvansky et al. (1993) when the material involved people and locations. Thus, we do not regard our results as contradicting theirs.

Table 4
Mean Latencies, Error Rates, and Latency Predictions for 2 and 4 Fans

Condition	Fan 2	Fan 4
Targets		
Facilitation		
Mean latency (in s)	1.33	1.42
Latency prediction	1.36	1.41
Error rate (%)	6.2	6.0
Interference		
Mean latency (in s)	1.57	1.58
Latency prediction	1.53	1.62
Error rate (%)	10.5	14.0
Suppression		
Mean latency (in s)	1.43	1.51
Latency prediction	1.43	1.49
Error rate (%)	4.9	7.8
Control		
Mean latency (in s)	1.42	1.53
Latency prediction	1.43	1.49
Error rate (%)	5.5	8.9
Foils		
High		
Mean latency (in s)	1.51	1.60
Latency prediction	1.55	1.60
Error rate (%)	7.5	9.9
Mixed		
Mean latency (in s)	1.58	1.67
Latency prediction	1.59	1.64
Error rate (%)	8.3	13.8
Low		
Mean latency (in s)	1.59	1.70
Latency prediction	1.62	1.68
Error rate (%)	11.7	14.4

The second and third observations concern the effect of probe type in Table 4. The second observation is that in neither the latency nor the error data was there any evidence for a suppression effect like that obtained by M. C. Anderson and Spellman (1995) with their materials. In contrast, there was statistically strong evidence in the targets for both a facilitation and an interference effect, indicating that our failure to find a suppression effect was not due to lack of power. The third observation is that, as predicted, there was an effect of target frequency on foil rejection. Foils are more rapidly rejected when there is a high-frequency target that involves one of the concepts. This supports the hypothesis that foils are rejected by retrieving a partially matching target.

Clearly, we failed to find an analog of M. C. Anderson and Spellman's (1995) suppression result in our paradigm. As we noted earlier, there are many procedural differences between the typical fan paradigm and M. C. Anderson and Spellman's paradigm. Thus, we do not view our research as challenging theirs. Rather, we view this research as showing that suppression cannot be the cause of the fan effect.

Mathematical Model

We fit the ACT–R model to the target and foil latencies. It is essentially the same model as described in the introduction, with parameters I for intercept, F for latency scale, and S for the associative strength constant. Recall that the formula for estimating a specific strength of connection between a concept j and a fact i is

$$S_{ji} = S + \ln [P(i|j)], \quad (2)$$

where $P(i|j)$ is the probability that i will have to be retrieved when j is present in the probe (based on the past proportion of occurrences where this was true). Up until now, we have simply set this as $1/f_j$, where f_j is the fan out of concept j. This was justified on the basis that all the facts about a concept j were presented equally often. In this experiment, the facts occurred with different frequencies during the RT test phase. In this experiment, facts could occur with probability five sixths about a 2-fan concept if the fact were in the facilitation condition, with probability one half if it were in the control condition, and with probability one sixth if it were in the interference condition. Similarly, facts could occur with probabilities five twelfths, one fourth, or one twelfth about 4-fan concepts. These were the probabilities we used for the $P(i|j)$ in Equation 2.

We fit the model to minimize the squared deviations between its predictions and the data in Table 4. We estimated I, the intercept parameter, to be 1,197 ms; F, the scale parameter, to be 773 ms; and S, the associative strength parameter, to be 2.50 ms. These parameters are similar to the estimates in the prior fits. The overall correlation between prediction and observation was .956. The average mean deviation of prediction from observation was 29 ms, which is good considering that the standard deviation of the means was 31 ms (estimated from the Participant × Condition interactions for targets and foils, ignoring the factors of blocks and which terms had their fan manipulated). Next, we work through two representative examples of how the predictions in Table 4 were derived.

1. *The 4-fan interference target:* There were two kinds of 4-fan interference targets. One example of such an item is *soldier–tunnel* from Table 3. This consisted of a 4-fan concept *tunnel* and a 2-fan concept *soldier.* According to Equation 2, the strength of association from the 4-fan concept *tunnel* to the fact *soldier–tunnel* is $S + \ln (1/12) = 0.02$. The 2-fan element, *soldier,* did not suffer from interference from the high-frequency fact *cowboy–tunnel.* Thus, the probability of *soldier–tunnel* given *soldier* remained one half throughout the experiment, reflecting the 2 fan. This means that the strength from *soldier* to the *soldier–tunnel* fact is $S + \ln (1/2) = 1.81$. Then, the expected activation given Equation 5 and our simplification that $B' = 0$ is $A = .33 * .02 + .33 * 1.81 = 0.61$. The other type of 4-fan interference fact from Table 3 is *cowboy–cage,* where the 4-fan element *cage* did not suffer interference from the high-frequency *cowboy–tunnel,* but the 2-fan element *cage* did. Similar mathematics reveal that it has the same expected activation of A = 0.61. Thus, our prediction based on Equation 3 is

$$\text{time} = 1{,}197 + 773(e^{-.61}) = 1{,}618 \text{ ms}.$$

2. *The 4-fan high foil:* An example of such a foil is *fireman–tunnel* from Table 3. The prediction in this case depended on whether the participant retrieved from the 4-fan concept *tunnel* or the 2-fan concept *fireman.* In either case, we assumed that a high-frequency target *cowboy–tunnel* (or whatever frequent fact had been studied with *fireman*) would be retrieved and used as a basis for rejection. If the participant retrieved from the 4-fan concept *tunnel,* the associative strength would be $S + \ln (5/12) = 1.62$, and activation (based on Equation 5′) would be $A = .333 * 1.62 = 0.54$. If the participant retrieved from the two-fan concept *fireman,* the associative strength would be $S + \ln (5/6) = 2.32$, and activation would be $A = .333 * 2.32 = 0.77$. Assuming an equal mixture of both retrievals, latency is predicted to be

$$\text{time} = 1{,}197 + 773\left(\frac{e^{-.54} + e^{-.77}}{2}\right) = 1{,}601 \text{ ms}.$$

Conclusions

The ACT–R theory of the fan effect has been able to account for the original data of J. R. Anderson (1974), recent data thought to be contradictory to the ACT theory (Radvansky et al., 1993), and results from a new experiment. There are three theoretical ideas that are important to this account. The most fundamental concept is that the strength of the association between a concept and a fact in memory is adjusted during prior experiences to reflect the statistical regularity with which that concept predicts that fact. As the fan increases, a concept becomes a poorer predictor, on average. However, as the results in Table 4 show, this can be the reverse for specific facts. Participants are slower for

2-fan interference facts than they are for 4-fan control or facilitation facts. This is because the probability of an interference fact occurring in the presence of a 2-fan concept is lower than the probability of a control or facilitation fact occurring in the presence of a 4-fan concept. Thus, the critical variable is probability not fan.

In addition to this basic insight, which ACT–R shares with the earlier ACT*, there are two new ideas that ACT–R contributes to help explain results in the literature. One is that it is possible to vary the amount of weighting given to various types of concepts. The larger the weighting given to one concept relative to another, the larger is the expected effect of that concept's fan. We have evidence suggesting that participants may weigh more heavily concepts that are more concrete and that tend to be better cues to memory. More generally, ACT–R predicts larger fan effects for concepts that receive greater attention.

The other contribution of the ACT–R model concerns the process of foil rejection. The basic idea is that foils are rejected by trying to retrieve a fact. If the retrieved fact does not match the probe, participants respond negatively. Because facts that mismatch the probe receive less activation than facts that match the probe, and because facts are retrieved in the order of their level of activation, there is little danger of falsely rejecting a target because of this mechanism. This offers a way to unify the theory of fan effects for targets and foils. It also succeeds in predicting the mean latency difference between targets and foils. Thus, there is finally a mechanistic theory of foil rejection that simultaneously accounts for the relative size of the fan effects for targets and foils as well as their mean latencies.

References

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

Anderson, J. R. (1976). *Language, memory, and thought.* Hillsdale, NJ: Erlbaum.

Anderson, J. R. (1981). Effects of prior knowledge on memory for new information. *Memory & Cognition, 9,* 237–246.

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

Anderson, J. R. (1990). *Cognitive psychology and its implications* (3rd ed.). New York: Freeman.

Anderson, J. R. (1993). *Rules of the mind.* Hillsdale, NJ: Erlbaum.

Anderson, J. R., & Bower, G. H. (1973). *Human associative memory.* New York: Holt, Rinehart & Winston.

Anderson, J. R., & Lebiere, C. (Eds.). (1998). *The atomic components of thought.* Mahwah, NJ: Erlbaum.

Anderson, J. R., & Paulson, R. (1978). Interference in memory for pictorial information. *Cognitive Psychology, 9,* 178–202.

Anderson, J. R., Reder, L. M., & Lebiere, C. (1996). Working memory: Activation limitations on retrieval. *Cognitive Psychology, 30,* 221–256.

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

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

Cantor, J., & Engle, R. W. (1993). Working-memory capacity as long-term memory activation: An individual-differences approach. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 1101–1114.

Conway, A. R. A., & Engle, R. W. (1994). Working memory and retrieval: A resource-dependent inhibition model. *Journal of Experimental Psychology: General, 123,* 354–373.

Estes, W. K. (1950). Toward a statistical theory of learning. *Psychological Review, 57,* 94–107, 110–115, 375.

Keenan, J. M., & Baillet, S. D. (1980). Memory for personally and socially significant events. In R. S. Nickerson (Ed.), *Attention and performance* (Vol. 8, pp. 651–670). Hillsdale, NJ: Erlbaum.

Kimberg, D. Y. (1994). *Executive functions, working memory, and frontal lobe function.* Unpublished doctoral dissertation, Carnegie Mellon University.

King, D. R. W., & Anderson, J. R. (1976). Long-term memory search: An intersecting activation process. *Journal of Verbal Learning and Verbal Behavior, 15,* 587–605.

Lewis, C. H., & Anderson, J. R. (1976). Interference with real world knowledge. *Cognitive Psychology, 8,* 311–335.

Myers, J. L., O'Brien, E. J., Balota, D. A., & Toyofuku, M. L. (1984). Memory search without interference: The role of integration. *Cognitive Psychology, 16,* 217–242.

Paivio, A. (1971). *Imagery and verbal processes.* New York: Holt, Rinehart & Winston.

Peterson, S. B., & Potts, G. R. (1982). Global and specific components of information integration. *Journal of Verbal Learning and Verbal Behavior, 21,* 403–420.

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

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

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

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

Reder, L. M., & Ross, B. H. (1983). Integrated knowledge in different tasks: The role of retrieval strategy on fan effects. *Journal of Experimental Psychology: Learning, Memory and Cognition, 9,* 55–72.

Reder, L. M., & Wible, C. (1984). Strategy use in question-answering: Memory strength and task constraints on fan effects. *Memory & Cognition, 12,* 411–419.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. Black & W. F. Prokasy (Eds.), *Classical conditioning: II. Current research and theory* (pp. 69–99). New York: Appleton-Century-Crofts.

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

Zbrodoff, N. J. (1995). Why is 9 + 7 harder than 2 + 3? Strength and interference as explanations of the problem-size effect. *Memory & Cognition, 23,* 689–700.

Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin, 123,* 162–185.

Received December 11, 1996
Revision received July 22, 1997
Accepted November 10, 1997 ■