On the Automatic Activation of Attitudes

Russell H. Fazio, David M. Sanbonmatsu, Martha C. Powell, and Frank R. Kardes
Indiana University

We hypothesized that attitudes characterized by a strong association between the attitude object and an evaluation of that object are capable of being activated from memory automatically upon mere presentation of the attitude object. We used a priming procedure to examine the extent to which the mere presentation of an attitude object would facilitate the latency with which subjects could indicate whether a subsequently presented target adjective had a positive or a negative connotation. Across three experiments, facilitation was observed on trials involving evaluatively congruent primes (attitude objects) and targets, provided that the attitude object possessed a strong evaluative association. In Experiments 1 and 2, preexperimentally strong and weak associations were identified via a measurement procedure. In Experiment 3, the strength of the object–evaluation association was manipulated. The results indicated that attitudes can be automatically activated and that the strength of the object–evaluation association determines the likelihood of such automatic activation. The implications of these findings for a variety of issues regarding attitudes—including their functional value, stability, effects on later behavior, and measurement—are discussed.

Our focus in this article is on the activation of attitudes from memory. The essential question to be addressed is whether attitudes are capable of being activated automatically upon the individual's encountering the attitude object. Consider such an encounter. One possibility is that the individual's attitude will be activated spontaneously and without any conscious effort on his or her part upon observation of the attitude object. On the other hand, it might be that activation of the attitude requires that the individual engage in a far more reflective process in which he or she actively considers his or her attitude toward the object. Our concern is with the extent to which the former possibility occurs and the degree to which the likelihood of its occurrence depends upon characteristics of the attitude in question.

The two possibilities regarding attitude activation outlined above correspond to the distinction offered by cognitive psychologists between automatic and controlled processes (e.g., Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1977). Shiffrin and Dumais (1981) characterized as automatic any process that leads to the activation of some concept or response "whenever a given set of external initiating stimuli are presented, regardless of a subject's attempt to ignore or bypass the distraction" (p. 117). The key feature of such automatic activation, then, is its inescapability. The implication for attitudes is that, upon presentation of an attitude object, an individual's attitude would be activated despite the lack of any reflection whatsoever on his or her part. In contrast, a controlled process requires the active attention of the individual. Upon becoming aware of a situational cue implying the importance of considering one's attitude toward an object, the individual might attempt to retrieve a previously stored evaluation of the attitude object or might actively construct such an attitude on the spot. In either case, the process is reflective and active in nature.

The occurrence of an automatic process requires the existence of a previously well-learned set of associations or responses. For example, Shiffrin and Schneider (1977) observed that a target stimulus developed the ability to attract attention automatically only following extensive training. The experimenters first trained subjects to respond to a set of characters (letters). On the critical trials of a subsequent task, these characters served as distractor items. That is, one such character might appear on a display in a location that was irrelevant to the subject's task. Despite this irrelevance, these characters to which the subject had earlier been trained to attend did attract attention, as indicated by relatively poorer performance on the primary task. Shiffrin (in press) reviewed a number of such investigations concerning the development of automatism.

Given that automatic processes require such well-learned responses, it appears doubtful that automatic activation is likely for all of the attitudes that an individual might hold. Only for well-learned ones is the expectation of automatic activation even a possibility. Social psychologists have long recognized that attitudes vary in their "strength". Indeed, a variety of attempts have been made to quantify and assess the centrality or importance of an attitude issue for a given individual. The notion of ego-involvement in the context of social judgment theory serves as an illustration of such an approach (Hovland, Harvey, & Sherif, 1957; M. Sherif & Cantril, 1947). More recently, various indices

The present research was supported by Grant MH 38832 from the National Institute of Mental Health.

The authors thank Paget Gross, Margaret Intons-Peterson, Steven Sherman, and Richard Shiffrin for their helpful comments on an earlier draft of the article.

Correspondence concerning this article should be addressed to Russell H. Fazio, Department of Psychology, Indiana University, Bloomington, Indiana 47405.

Journal of Personality and Social Psychology, 1986, Vol. 50, No. 2, 229–238
 0022-3514/86/$00.75

of the "strength" of an attitude have been identified as moderators of the relation between attitudes and behavior. As examples, the confidence with which an attitude is held (Sample & Warland, 1973; Fazio & Zanna, 1978a, 1978b); how clearly defined the attitude is, as measured by the width of the latitude of rejection (C. Sherif, Kelly, Rodgers, Sarup, & Tittler, 1973; Fazio & Zanna, 1978a); and the consistency between affective and cognitive components of the attitude (Norman, 1975) have each been found to relate to attitude–behavior consistency.

Relevant to this idea of attitudes varying in strength is the so-called *attitude/nonattitude* distinction. A number of years ago, both Hovland (1959) and Converse (1970) attempted to reconcile differences that had been observed between survey and laboratory research on attitude change. In so doing, they each—but Converse in particular—focused on a distinction between attitudes and nonattitudes. The distinction centered on the observation that a person may respond to an item on an attitude survey even though that particular attitude does not really exist in any a priori fashion for the individual. The attitude object may be one that the individual has not even considered prior to administration of the attitude survey. For Converse (1970), the attitude/nonattitude distinction centered on measurement error. An individual's nonattitude was characterized by unreliable measurement (in fact, virtually random responding) across the waves of a panel survey.

The attitude/nonattitude dichotomy might be more fruitfully conceived as a continuum. At one end of the continuum is the nonattitude. No a priori evaluation of the attitude object exists. As we move along the continuum, an evaluation does exist and its accessibility from memory grows increasingly strong. At the other extreme of the continuum, then, is a well-learned attitude that is highly accessible from memory.

A particular conception of attitudes underlies this view of the attitude/nonattitude continuum. In a recent series of experiments concerning attitude accessibility, Fazio and his colleagues (Fazio, Chen, McDonel & Sherman, 1982; Fazio, Powell, & Herr, 1983; Powell & Fazio, 1984) proposed that attitudes be viewed as simple associations between a given object and a given evaluation. The term *object* is used in a broad sense. Individuals may have evaluations of a wide variety of potential attitude objects, including social issues, categories of situations, categories of people, and specific individuals, as well as physical objects. Likewise, the term *evaluation* is used in a broad sense. It may range in nature from a very "hot" affect (the attitude object being associated with a strong emotional response) to a "colder" more cognitively-based judgment of one's affect (feelings of favorability or unfavorability) toward the object (see Abelson, Kinder, Peters, & Fiske, 1982; Zanna & Rempel, 1984).

More relevant to the present purposes is the notion of association between the attitude object and the evaluation, regardless of the precise nature of this evaluation. The definition of an attitude implies that the strength of an attitude, like any other construct based on associative learning, can vary. That is, the strength of the association between the object and the evaluation can vary. This associative strength may determine the accessibility of the attitude from memory and the likelihood that the attitude will be activated automatically upon the individual's encountering the attitude object. Only if it is strongly associated with the attitude object is it likely that the evaluation will be spontaneously activated upon mere presentation of the attitude object. In general, what is being suggested is that the activation of one's affect toward an object (be it a hot or a cold linkage) depends on the strength of the association.

In testing this view of attitudes as object–evaluation associations, Fazio and his associates (Fazio et al., 1982; Powell & Fazio, 1984) employed latency of response to an attitudinal inquiry as a measure of the associative strength. Subjects were asked to indicate as quickly as possible whether they felt positively or negatively toward a given attitude object. Subjects who had been induced to express their attitudes repeatedly—which should have the consequence of strengthening the object–evaluation association—were able to respond relatively quickly to these direct inquiries about their attitudes. For example, Powell and Fazio (1984) manipulated the number of times that an attitude was expressed in a within-subjects design by varying the number of semantic differential scale items that appeared relevant to a given attitude issue. In this way, subjects expressed their attitudes zero, one, three, or six times toward a given object.[1] In a subsequent task, subjects were presented with each attitude object and were instructed to make a good or bad judgment about each object as quickly as possible. Latency of response (from stimulus onset to response) was found to relate to the number of previous attitudinal expressions. The greater the number of expressions, the faster the latency of response to the attitudinal inquiry.

These findings imply that attitudes characterized by strong object–evaluation associations may be more accessible from memory. However, it is important to recognize that these findings are not at all informative with respect to the issue of whether the attitude activation stems from an automatic or a controlled process. Responding quickly to a direct attitudinal inquiry does not necessarily mean that the stored evaluation was activated automatically. Instead, the evaluation simply may have been retrieved efficiently via an effortful, controlled process.

In order to examine the automatic activation of attitudes from memory, the present research employed a priming procedure. The procedure is a variant of a now well-tested method commonly employed to investigate automatic processing. It involves consideration of the extent to which the presentation of a prime automatically activates concepts that facilitate responding to a target word. For example, Neely (1977) found that presentation of a category label as a prime (e.g., bird) facilitated the speed with which subjects could identify a subsequently presented target word as a word, provided that the target was semantically related to the category (e.g., robin). The technique has been used to study activation from memory in a variety of contexts, including text processing (e.g., McKoon & Ratcliff, 1980; Ratcliff & McKoon, 1978) and spatial representations (e.g., McNamara, Ratcliff, & McKoon, 1984), as well as semantic relations (e.g., Neely, 1976, 1977; deGroot, 1983).

In the present context, the subjects' primary task was to indicate as quickly as possible whether a target adjective (e.g., pleasant) had a positive or negative connotation. Latency of response served as the dependent measure. Our concern was with

[1] This manipulation had no effect on the extremity of the final attitudinal expression. Thus, the results appear to be due to the strengthening of the object–evaluation association as a consequence of repeated attitudinal expression.

the extent to which such a judgment would be facilitated by the presentation of an attitude object as the prime. We reasoned that presentation of an attitude object would automatically activate any strong association to that object. Such activation is assumed to spread along the paths of the memory network, including any evaluative associations. Consequently, the activation levels of associated evaluations are temporarily increased. If a target word that corresponds in valence to one of these previously activated evaluations is subsequently presented for judgment, then less additional activation is required for the activation level of the target word to reach threshold and, consequently, for a judgment to be made. Responding to a target word that has received some activation as a result of presentation of the prime is thus facilitated. That is, the individual should be able to respond relatively quickly.

As an example, let's assume that the attitude object *vodka* is evaluated positively by an individual. Presentation of vodka as the prime may automatically activate a positive evaluation. If the target adjective that is presented is also positive, then the individual may be able to indicate relatively quickly that the target has a positive connotation. That is, facilitation should occur. In a similar manner, facilitation is expected in the case of a negatively valued object serving as the prime when it is followed by a negative target adjective, as in *cockroach/disgusting*. What is meant by facilitation is simply that the latency is faster in such cases than in a trial involving the same target word preceded by a letter string (e.g., BBB). Such trials provide a no-prime baseline. Thus, the technique relies on the presence of facilitation as an indication that the evaluation associated with the primed attitude object has been activated upon its mere observation.

Of course, we would not expect such facilitation to occur for all attitude objects. Relating the methodology to the previous discussion of the attitude/nonattitude continuum, we would expect facilitation to occur only if the object–evaluation association is quite strong. In Experiments 1 and 2, the strength of the object–evaluation association was assessed with regard to a large number of potential attitude objects and strong versus weak primes were selected on an individual basis for each and every subject. In Experiment 3, the strength of the object–evaluation association was manipulated.

Experiment 1

The first experiment involved selecting attitude objects toward which a given individual possessed a strong versus a weak evaluative association and then testing whether those objects produced facilitation when presented as primes in the major experimental task. In order to assess the strength of object–evaluation associations, we employed the operationalization that had been used successfully in the studies of attitude accessibility mentioned earlier (Fazio et al., 1982; Powell & Fazio, 1984). Recall that the findings from this previous research indicated that latency of response to an attitudinal inquiry appears to index the strength of an object–evaluation association satisfactorily.

In the present experiment, attitude objects toward which a given subject displayed very fast or very slow latencies of response to an attitudinal inquiry were identified. If the latency was fast, we consider the object–evaluation association strong and, hence, facilitation should occur in the procedure described earlier. That is, positive target adjectives should be identified as having a positive connotation relatively more quickly when preceded by a positively valued object. Likewise, negative target adjectives should be identified as having a negative connotation relatively more quickly when preceded by a negatively valued object. Such facilitation is far less likely in the case of a weak object–evaluation association, as indicated by a slow latency of response to the direct inquiry. Thus, the hypothesis leads to a prediction of a three-way interaction (Strength of Association × Valence of Prime × Valence of Target). Greater facilitation is expected on trials involving congruent valences than on trials involving incongruent valences (i.e., a simple interaction of Prime Valence × Target Valence) for primes involving a strong evaluative association but not for primes involving a weak association.

Method

Subjects. Twenty-two Indiana University undergraduates participated in the experiment in partial fulfillment of an introductory psychology course requirement.

Procedure. Subjects were told that the experiment concerned word recognition and meaning and that a number of different tasks relevant to word judgment would be performed during the course of the experiment. They were also told that these tasks would grow increasingly complex as we progressed through the procedure.

The experimental procedure consisted of two major phases, the first devoted to prime selection and the second involving the actual priming task. A list of 70 attitude objects (including the names of some individuals, animals, foods, social groups, nations, activities, and physical objects) formed the pool of potential primes. Subjects were told that the first and simplest word-judgment task that they would be performing involved the presentation of a single word on the computer screen on any given trial. Their task was to press a key labeled *good* or a key labeled *bad* as quickly as possible to indicate their judgment of the object. Subjects were instructed to maximize both the speed and accuracy of their responses. The presentation was controlled by an Apple II+ computer. The order in which the words were presented was randomized for each subject. A given word remained visible on the screen until the subject responded. A 3-s interval separated each trial. The subject's response was recorded, along with the latency of response (from word onset to response) to the nearest millisecond. Subjects' performance of this task was preceded by a block of practice trials involving different words than those used as the potential primes, so as to familiarize subjects with the procedure.

After performing the task, subjects were excused from the laboratory for a short break. During this time, 16 words were selected on the basis of the subject's data as the primes. Four words were selected in each of four categories: strong good, strong bad, weak good, and weak bad. The 4 words toward which the subject had responded *good* and the 4 toward which the subject had responded *bad* most quickly served as the strong primes. The 4 good and the 4 bad words involving the slowest latencies served as the weak primes.[2] These 16 words, along with four different

[2] Which attitude objects from the pool of 70 potential primes were selected for use as primes was quite idiosyncratic across subjects. Nevertheless, the following tabulation is intended to provide the reader with some sense of the nature of the attitude objects that served as priming stimuli. The most frequently selected objects in each of the four prime categories and the number of subjects for whom each object was selected are listed: strong good—gift (7), music (7), party (6), and cake (5); strong bad—death (11), hell (7), guns (6), and crime (5); weak good—crosswords (8), Republicans (7), Democrats (7), and rum (6); and weak bad—mazes (9), radiation (7), Democrats (7), and recession (6). Some of the other

strings of three identical letters (e.g., BBB), which were intended to provide nonprime baselines, were employed as the primes in the next task.

A list of 10 evaluative adjectives that were clearly positive in connotation (e.g., "appealing," "delightful") and a list of 10 adjectives that were clearly negative (e.g., "repulsive," "awful") were prepared. These words served as the target words in the next phase of the experiment. Subjects were told that this task was a more complex one involving their again making a judgment of a word, but that this time they would have to remember another word while making the judgment. They were informed that a memory word would be presented followed by an adjective. They were to press the good or bad key as quickly as possible to indicate whether the adjective had a positive or negative connotation and to then recite the memory word aloud. Subjects were told to recite the memory word, that is the prime, solely to ensure that they attended to the prime. (In the case of a letter string such as BBB, subjects were instructed to recite "Triple B".) A cassette recorder was positioned adjacent to the computer to bolster the presumption that the experimenter was concerned about the subjects' recitation of the memory word.

On any given trial, a prime was presented for 200 ms, followed by a 100-ms interval before onset of the target word. Thus, the interval between prime onset and target onset, commonly referred to as the stimulus onset asynchrony (SOA), was 300 ms. The target word disappeared upon the subject's pressing a key. A 4-s interval passed before presentation of the next prime.

A total of five blocks of trials were presented. Each block consisted of 20 trials, in which each of the 20 primes (including the four letter strings) and each of the 20 target adjectives were presented once. Within each of the five prime categories (strong good, strong bad, weak good, weak bad, and letter string), 2 of the primes were followed by positive adjectives and 2 by negative adjectives. Across blocks, each target adjective was paired once with a prime from each of the five prime categories (strong good, strong bad, weak good, weak bad, and letter string). Thus, a target adjective appeared equally often in each of the five prime conditions. As with the prime selection task, subjects underwent a series of practice trials before performing the actual task so as to familiarize them with the procedure.

Results and Discussion

Subjects committed very few errors in making judgments of the connotation of the target adjectives. The average error rate across subjects was 1.95%. In these few cases, the latency was excluded from the analysis. For each subject, the mean response latency in each of the 10 cells of the design (Five Prime Categories × Positive vs. Negative Targets) was computed. Facilitation scores were then computed. Each mean in a positive target condition was subtracted from the nonprime baseline provided by trials in which positive targets were preceded by a letter string. The same was done with respect to the negative target conditions.[3] The resulting facilitation scores are depicted in Figure 1.

A 2 (strength of association) × 2 (prime valence) × 2 (target valence) analysis of variance was performed on the facilitation scores. The analysis revealed that the expected three-way interaction was statistically significant, $F(1, 21) = 6.86$, $p < .02$. In the case of primes involving a strong evaluative association, the predicted interaction between prime valence and target valence was very apparent, $F(1, 21) = 15.30$, $p < .001$. Just as we predicted, facilitation occurred in the cases of congruency between

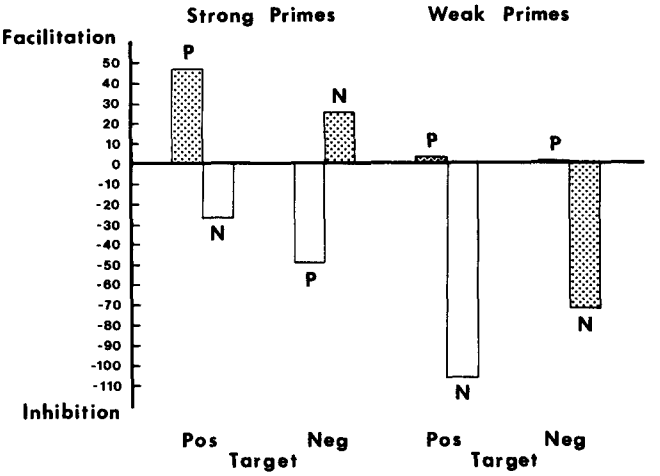

Figure 1. Mean facilitation scores in Experiment 1 as a function of the strength of the object–evaluation association, the valence of the prime (P = positive; N = negative), and the valence of the target adjective (Pos = positive; Neg = negative). (Conditions involving congruent valences are darkened; incongruent ones are not.)

the valence of the primed object and the valence of the target adjective, but not in the cases of incongruency. That is, facilitation is apparent for positively valued attitude objects when followed by a positive target and for negatively valued objects when followed by a negative target. Inhibition is apparent for positively valued attitude objects when followed by a negative target and for negatively valued objects when followed by a positive target.

As indicated by the significant three-way interaction, the data pattern is quite different for those objects involving a weak evaluative association. Most important, there is no interaction between prime valence and target valence, $F < 1$, and no evidence of facilitation in the case of the weak primes. The only effect is a main effect of the valence of the primed object, $F(1, 21) = 14.58$, $p < .001$. Negative objects produced inhibition regardless of the valence of the target adjective. Apparently, the negative objects involving a weak association somehow distracted subjects from the central task. However, this effect should be interpreted cautiously because, unlike the effect observed for strong primes, it does not replicate in the later experiments to be reported.

An important point does need to be made about facilitation versus inhibition in the present procedure. It should be noted that strong primes produced just as much inhibition with regard to targets of incongruent valence as they produced facilitation of targets of congruent valence. The explanation may center on the letter-string trials that were employed as a nonprime baseline. There is considerable discussion in the cognitive literature about the nature of stimuli appropriate for use in baseline trials (see deGroot, Thomassen, & Hudson, 1982; Jonides & Mack, 1984). Ideally, such stimuli should have no potential for inducing spreading activation, but should be identical to the actual primes in all other respects, including physical appearance, potential to

objects that also were selected for use as primes relatively frequently, but in varying categories, included vodka (8), snow (8), spider (8), television (7), dentist (6), sports (6), storms (6), Reagan (5), coffee (5), and Iran (4).

[3] For the benefit of any reader interested in considering raw latencies, the positive target and negative target baselines were 1,065 ms and 1,090 ms, respectively.

alert subjects, ease of encoding, and memory demands (Jonides & Mack, 1984). Only with the ideal baseline is it appropriate to consider a prime as having produced *facilitation* or *inhibition* in any absolute sense. In recognition of this difficulty, we use these terms in a relative sense throughout this article. Facilitation refers to faster responding to a target word when it is preceded by a prime than when it is preceded by a letter string and inhibition refers to slower responding. Whether the letter-string trials provide an estimate of the true baseline is uncertain. Thus, the zero point of our facilitation–inhibition scale may be a misestimate. What appears to be relative inhibition may actually represent slight facilitation in an absolute sense.

Indeed, we suspect that the letter-string trials may overestimate the true baseline and, consequently, underestimate the amount of absolute facilitation that is occurring. This suspicion is based on the fact that subjects were required to recite the prime aloud at the end of each trial and on the possibility that the letter strings represented less of a memory load than did the attitude object primes. This lesser memory load would allow for responding that is faster than the ideal nonprime, one that is equivalent to the attitude object primes in terms of required memory load.

Regardless of whether a given effect is to be labeled as facilitation or inhibition, the present findings clearly indicate that individuals were able to respond to target adjectives more quickly when the targets were preceded by attitude objects congruent in valence than when they were preceded by attitude objects of incongruent valence. However, as expected, this held true only for attitude objects toward which a strong evaluative association existed.

On the basis of these findings, it appears that at least some attitudes may be activated from memory automatically upon mere presentation of the attitude object. What is critical is that the present evidence for subjects' attitudes having been activated is found in a situation in which the subject was merely exposed to the attitude object. The subjects were never asked during the second phase of the experiment to consider their attitudes. Nor was it to the subjects' advantage to do so, for the subjects' task was simply to respond to the target word and then to recite the memory word. Nevertheless, despite this irrelevance of attitudes to the immediate task concerns, exposure to objects for which subjects presumably possessed strong affective associations appears to have prompted activation of the associated evaluation.

Thus, the very nature of the task leads to the suggestion that the facilitation observed in the case of the strong primes was a result of automatic, rather than controlled, processing. Nevertheless, it might be argued that subjects did for some reason actively consider their evaluations of the prime and, hence, were "prepared" for a target word of congruent valence. To explain the findings, such an interpretation would have to maintain that the SOA that was employed allowed sufficient time for subjects to actively retrieve their evaluation in the case of strong primes, but was insufficient for such active retrieval in the case of weak primes. In order to examine this possibility, a second experiment involving manipulation of the SOA was conducted.

Experiment 2

If the facilitation observed in Experiment 1 were due to a controlled process, then allowing the subjects more time should, if anything, enhance the extent of facilitation. Most important, facilitation might be observed even in the case of weak primes. On the other hand, if the task is such, as we have argued, that the findings in Experiment 1 reflect automatic processing, then no such facilitation is to be expected for weak primes even at a longer SOA. Assuming that the interpretation in terms of automatism is valid, then whether facilitation in the case of strong primes is observed at a longer SOA will depend on the level of activation of the associated evaluation at the time the target word is presented. If the level of activation has dissipated (possibly due to its irrelevance to the immediate task concerns) then no facilitation is to be expected at the longer SOA. If the level of activation has not yet returned to baseline, then some facilitation might be expected.

Method

Subjects. Twenty-three individuals who had responded to a newspaper advertisement participated in the experiment in return for a payment of $6.

Procedure. The experimental procedure followed that employed in Experiment 1. The only major difference was that, following prime selection, subjects underwent the actual priming task twice. They did so once with an SOA of 300 ms and once with an SOA of 1000 ms. The order in which they did so was counterbalanced across subjects.

The only other procedural changes made were minor ones aimed at enhancing the power of the experiment. The pool of potential primes employed in the prime selection phase of the experiment was expanded to 92 attitude objects.[4] In addition, a few target adjectives that had produced relatively short or long latencies when preceded by letter strings were replaced by other words.

Results and Discussion

As in Experiment 1, errors were minimal (mean error rate = 1.39%) and when they did occur, the respective latency was omitted from the analysis. Facilitation scores[5] are depicted in Figure 2. Because the order of the blocks of trials involving SOAs of 300 versus 1,000 ms did not qualify any of the effects to be reported, the facilitation scores are presented collapsed across the order variable. As is clear from Figure 2, left column, the findings with an SOA of 300 ms replicate those observed in Experiment 1 when one considers attitude objects involving a strong evaluative association. Most important, there was a significant interaction of prime and target valence for strong primes at this SOA, $F(1, 22) = 4.87$, $p < .05$. Facilitation was greater in the case of congruent valences than in the case of incongruent va-

[4] The attitude objects most frequently selected for use as primes in each of the four categories (and the number of subjects for whom the object was selected) were as follows: strong good—music (7), friend (7), dancing (5), and cake (5); strong bad—war (8), death (8), cancer (6), and rats (6); weak good—Monday (9), dormitory (5), landlords (5), and Reagan (4); and weak bad—anchovies (8), landlords (6), exams (6), and recession (6). Some of the other objects that also were selected for use as primes relatively frequently, but in varying categories, included priest (8), fraternity (7), dentist (6), liver (6), mosquito (6), worms (6), guns (5), spider (5), taxes (5), and disco (4).

[5] The positive and negative target baselines were 830 ms and 880 ms, respectively, at the SOA of 300, and 770 ms and 790 ms, respectively, at the SOA of 1,000.

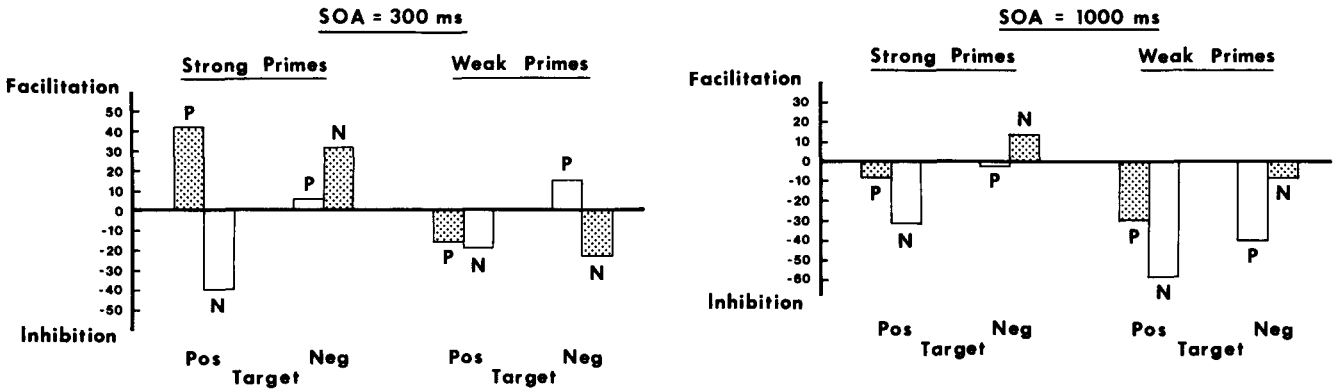

Figure 2. Mean facilitation scores in Experiment 2 as a function of SOA (stimulus onset asychrony), the strength of the object–evaluation association, the valence of the prime (P = positive; N = negative), and the valence of the target adjective (Pos = positive; Neg = negative). (Conditions involving congruent valences are darkened; incongruent ones are not.)

lences. Furthermore, the extent of facilitation in the congruent cases (M = 37 ms) differed significantly from zero, $t(22) = 2.17$, $p < .05$. No significant facilitation was apparent in any of the other cells of the design—strong primes presented at the longer SOA or weak primes at either SOA. This pattern of data led to the observation of a significant SOA × Strength of Association × Prime Valence × Target Valence interaction, $F(1, 22) = 4.35$, $p < .05$. Only in the case of strong primes and the shorter SOA was facilitation found on trials involving congruent evaluations.

Thus, even when subjects were provided with additional time, no evidence of facilitation was observed for weak primes. It does not appear that the evaluative association was activated upon presentation of the attitude object when the association was weak in nature. This null finding tends to undermine the plausibility of the controlled processing alternative outlined earlier. If subjects had been actively retrieving their attitudes, then greater facilitation was to be expected at the longer SOA.

Apparently, facilitation in the present task is the result of automatic activation of the evaluation upon presentation of the attitude object. However, such automatic activation requires the existence of a strong association between the attitude object and the evaluation. The findings regarding SOA imply that the level of activation in such cases dissipated quickly (or, conceivably, was actively suppressed). Such quick dissipation may have been a consequence of the irrelevance of the subject's attitudes to the major task that was occupying the subject's attention, that is, identifying the connotation of the target adjective. In effect, presentation of the target adjective 1,000 ms after presentation of the attitude object appears to have been too late for the prime to facilitate responding to adjectives of the same valence. At the 300-ms interval, on the other hand, the level of activation of the associated evaluation was apparently sufficient to facilitate responding to evaluatively congruent adjectives. This finding is reminiscent of one from Neely's (1977) investigation. Subjects had been instructed that the prime *bird* implied that a target word corresponding to a body part would be presented. Despite this instruction, responses to a specific bird exemplar (e.g., robin) as the target were facilitated at a short SOA. However, no such facilitation was observed at a longer SOA. In a fashion parallel to the present case, the level of activation of whatever exemplars had been activated by the presentation of the category prime apparently dissipated quickly.

Experiment 3

In each of the two experiments reported thus far, evidence supportive of the possibility of automatic activation of one's attitude upon mere exposure to the attitude object was found. Such automatic activation was restricted, however, to attitudes involving a strong object–evaluation association. Associative strength was measured via latency of response to an attitudinal inquiry. Attitudes that individuals could report relatively quickly were considered to involve strong object–evaluation associations. Such attitudes were compared to ones that required more time for subjects to report. These two classes of attitudes were found to differ with regard to the likelihood of presentation of the attitude object automatically activating the attitude from memory. We have argued that the strength of the object–evaluation association was the critical difference responsible for this differential likelihood. Nevertheless, as is the case any time a conceptual variable is measured rather than manipulated, other differences might exist between the two classes of attitudes that were identified via our measurement technique.

Experiment 3 was aimed at demonstrating more conclusively the critical importance of the strength of the object–evaluation association by manipulating rather than measuring it. Subjects were induced to express their attitudes toward a number of attitude objects repeatedly, as in the study described earlier by Powell and Fazio (1984). Another set of attitude objects was presented equally often but subjects were asked to make a nonevaluative judgment regarding each of these objects. In this way, the object–evaluation association was strengthened for some attitudes and not for others. These objects then served as primes in the subsequent task. The experimental design was the same as that of Experiment 2. The only difference was that strong versus weak primes were created experimentally rather than being selected on the basis of measured preexperimental strength. A

four-way interaction similar to that found in Experiment 2 is to be expected. That is, facilitation should be greater in the cases of congruency between the evaluation of the prime and the evaluation of the target than in the cases of incongruency only for strong primes presented 300 ms before the target words.

Method

Subjects. Eighteen Indiana University undergraduates participated in the experiment in partial fulfillment of an introductory psychology course requirement.

Procedure. As before, subjects were led to believe that the experiment concerned word recognition and that the experimenter was interested in the "speed and accuracy with which people could perform various word recognition tasks." In the initial task, a word and a question appeared simultaneously on the computer screen. In some instances, the question asked "One syllable word?" and the subject was to respond by pressing a yes or no key as quickly as possible. In other instances, the question was "Good or bad?" and subjects responded by pressing the appropriate key as quickly as possible. The words and accompanying question were presented in a random order for each subject. Three seconds separated each trial. Subjects underwent a series of practice trials before performing the actual task so as to ensure their understanding of the task.

A total of 16 attitude objects served as the stimuli that were subject to manipulation. The words were selected from the pool of 92 potential primes employed in Experiment 2 on the basis of the response and response latency data from the prime selection phase of Experiment 2. Two relevant criteria were employed to guide this process. First, objects that were endorsed as positive or negative with near unanimity across the subjects were selected. Second, of these objects that produced nearly uniform responses, the ones with the longest average response latencies across subjects were chosen. In this way, 8 positively valued ("aquarium," "baby," "cake," "chocolate," "eagle," "Friday," "parade," and "silk") and 8 negatively valued ("divorce," "hangover," "litter," "radiation," "recession," "toothache," "virus," and "weeds") attitude objects were selected.[6] These words were then randomly divided into two lists each consisting of 4 positive and 4 negative objects. For any given subject, one list was associated with the attitude expression question and one with the syllable identification question. Which specific list was used for which purpose was counterbalanced across subjects.

The actual task consisted of five blocks of 30 trials each. In any given block, the 16 words of interest appeared once. Thus, across blocks, subjects expressed their attitudes toward 8 of the objects 5 times[7] and answered the syllable question with respect to each of the 8 words from the other list 5 times. In this way, the number of presentations of the critical words was held constant. Filler words were included to bring the number of trials in each block to 30. Sixteen filler words were presented once throughout the series of blocks and another 18 filler words were presented a total of 3 times across all blocks. Half of the filler words were paired with the evaluation question and half with the syllable question.

The 4 positive and the 4 negative objects toward which evaluations had been expressed repeatedly in the above task are referred to as the *repeated expression* primes. The 4 positive and 4 negative objects that had been paired with the syllable question served as the control primes. The actual priming task proceeded as in Experiment 2. Following the priming task, subjects completed a questionnaire that asked them to evaluate each of the 50 attitude objects that had been presented in the first phase of the experiment (i.e., both the experimental and the filler words). These evaluations were made on a 7-point scale with endpoints labeled *very good* and *very bad.*

Results

As in the previous experiments, the number of errors that subjects made in judging the connotation of the target adjectives was minimal (mean error rate = 1.94%). Any latencies associated with a response error were omitted from the analysis.

The facilitation scores[8] are presented in Figure 3. An analysis of variance on these data revealed the predicted four-way interaction of SOA × Repeated Expression × Prime Valence × Target Valence, $F(1, 17) = 9.22$, $p < .01$. Similar to what had been observed for preexperimentally strong primes in Experiment 2, the critical simple interaction of prime and target valence was found for the experimentally created strong primes (repeated expression condition) when the shorter SOA was involved, $F(1, 17) = 8.86$, $p < .01$. In this case, greater facilitation was observed when positively valued primes were followed by positive targets and when negatively valued primes were followed by negative targets than when the prime-target pairs were incongruent in valence. Furthermore, the extent of facilitation observed on these congruent trials (M = 73.61 ms) was once again statistically reliable, $t(17) = 4.37$, $p < .001$. Statistically reliable facilitation (M = 43.19 ms) was also apparent on congruent trials with control primes when the trials involved the 300-ms SOA, $t(17) = 3.30$, $p < .005$.[9] However, in a manner that is consistent with the hypothesis, the extent of such facilitation was significantly greater in the repeated expression condition than in the control condition, $t(17) = 2.61$, $p < .02$. No facilitation on congruent trials was apparent with either the repeatedly expressed or the control primes when the trials involved the longer SOA.

Note that the present data regarding the consequences of repeated attitudinal expression are not confounded by attitude extremity. That is, repeated expression did not enhance the extremity of subjects' self-reported attitudes at the end of the experiment. Attitudes toward each of the 16 experimental objects were examined by comparing the attitude scores of those subjects who had repeatedly expressed their attitudes toward the given object to the attitude scores of those who had not. Only 1 of these 16 comparisons, well within what is to be expected by chance alone, revealed a more extreme attitude among the repeated expression subjects. Such a null effect of repeated expression on attitude extremity is consistent with past research that has employed the repeated attitudinal expression manipulation. Powell and Fazio (1984) did not detect any reliable differences in attitude extremity as a function of repeated expression across 12 attitude issues. Thus, the present findings regarding attitude activation following repeated expression can be attributed to the resulting strength of the object–evaluation association and not to the extremity of the associated evaluation.

---

[6] An additional 6 subjects who participated in the experiment were not included in the analysis because their evaluative judgments of the objects chosen as experimental stimuli were not in accordance with these expectations. Four of these subjects disagreed concerning one of the attitude objects (e.g., judging negatively what had been selected as a positively valued object) and 2 did not concur with respect to two attitude objects.

[7] As is to be expected, subjects were significantly faster at indicating their evaluations the fifth time that they did so (M = 910 ms) than the first time (M = 1,330 ms), $F(1, 17) = 73.95$, $p < .001$.

[8] The positive and negative target baselines were 780 ms and 840 ms, respectively, at the SOA of 300 ms, 710 ms, and 760 ms, respectively, at the SOA of 1,000.

[9] Due to the selection requirement concerning unanimity, the control attitude objects in Experiment 3 involved stronger evaluative associations than did the weak primes in Experiments 1 and 2.

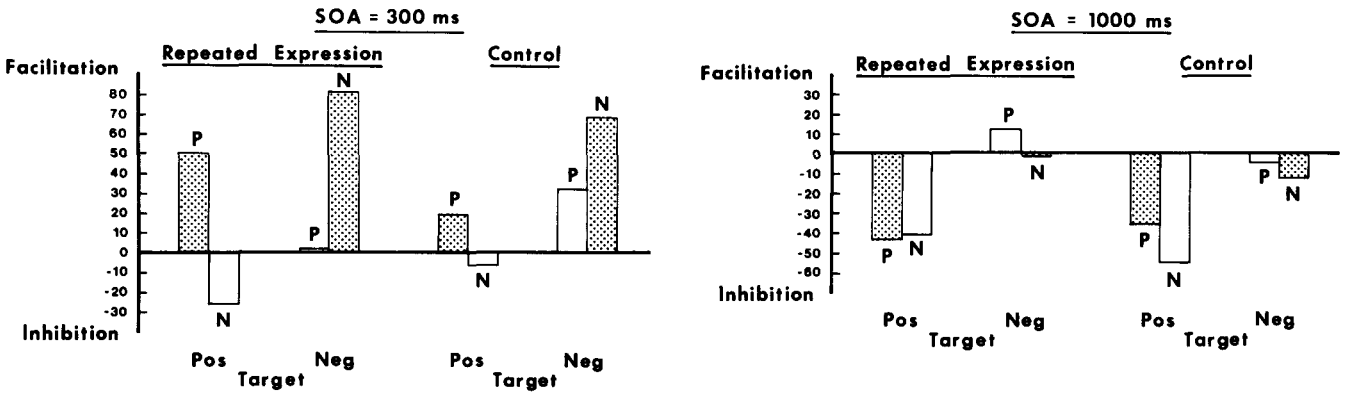

Figure 3. Mean facilitation scores in Experiment 3 as a function of SOA (stimulus onset asychrony), repeated expression, valence of the prime (P = positive; N = negative), and valence of the target adjective (Pos = positive; Neg = negative). (Conditions involving congruent valences are darkened; incongruent ones are not.)

General Discussion

Together, the three present experiments suggest that some sorts of attitudes indeed can be activated from memory automatically upon one's mere observation of the attitude object. Such activation appears to be both spontaneous and inescapable. Even though attitude activation was irrelevant to the task that subjects were required to perform, we found evidence that evaluations were activated upon exposure to appropriate attitude objects. However, the likelihood of such automatic activation of an attitude appears to depend upon the strength of the association between the attitude object and the evaluation. Regardless of whether this associative strength was measured or manipulated experimentally, evidence of automatic activation of the attitude was far stronger when the association could be characterized as strong.

It should be kept in mind that the view of attitude underlying the present research equates it with affect. Thus, the research essentially addresses how affect toward an object might be activated from memory. As mentioned earlier, individuals may have affect toward a wide variety of potential attitude objects, including social issues, categories of people, and specific individuals, as well as physical objects. The present findings, then, are of relevance to any broadly defined "object" toward which an individual possesses some affective linkage. As such, they indicate that affect can be activated automatically from memory in the same way that has been demonstrated previously for semantic knowledge. Just as a knowledge structure concerning some object may consist of bits of information organized in a network of associations to the object, so too may affect be linked to the object (cf. Fiske, 1982; Fiske & Pavelchak, in press). Furthermore, just as activation can spread from one node in the network to another (cf. Anderson & Pirolli, 1984; Ratcliff & McKoon, 1981), the present data indicate a spontaneous spreading of activation from the object to the affective association.

The findings also lend support to the utility of the proposed view of attitudes as evaluative associations with objects. These associations can vary in strength, ranging from *not existing at all* (the case of a nonattitude) to a *weak association* that is unlikely to be capable of automatic activation to a *strong association* that can be activated automatically. This continuum provides an interesting way of conceptualizing the strength of an attitude. The attitudes of two individuals with identical scores from some attitude measurement instrument may still differ markedly with regard to their strength, that is, with regard to their likelihood of activation upon the individual's mere exposure to the object. When they encounter the attitude object in a given situation, the attitude of one individual may be activated whereas the attitude of the other may not be. As a result, the two individuals may construe any information that is available concerning the object quite differently; that is, selective processing is more likely in the case of the individual whose attitude has been activated. A difference of this sort is apt to have a number of important implications.

First of all, the degree to which the attitude is likely to be evoked automatically is apt to affect the resistance of the attitude to counterinfluence. An attitude involving a strong association is apt to be activated upon the presentation of information concerning the attitude object and, as a result, color one's judgments of the information. Consistent with this notion, Wood (1982) found that attitude change in response to a persuasive communication is moderated by the degree to which individuals can rapidly retrieve from memory beliefs about the attitude object and past behaviors that they had performed relevant to the attitude object. The implication is that the stronger the object–evaluation association, the more resistant the attitude is to change and the greater the stability of the attitude over time.

A second implication concerns attitude–behavior consistency. If an attitude is activated automatically upon the individual's encountering the attitude object, it is far more likely to guide the individual's behavior toward that object than if it is not. Without activation of the attitude, behavior toward the object may proceed without the object having been considered in evaluative terms or on the basis of judgments of whatever features of the object happen to be salient in the immediate situation. In either case, the behavior may not be congruent with the individual's attitude.

Fazio (in press) has proposed a model of the process by which attitudes guide behavior that views activation of the attitude as critical in precisely this manner. The model assumes that behavior

in any given situation is largely a function of the individual's definition of the event that is occurring. The critical question concerns the extent to which the attitude will influence one's construction of the event. In some situations, a cue implying the relevance of attitudinal considerations may prompt individuals to access their attitudes from memory. However, in cue-free situations, it is the chronic accessibility of the attitude, that is, the strength of the object–evaluation association, that is important. If the attitude is activated upon one's observation of the attitude object, it is likely to lead to some selective processing of the information available in the immediate situation. Thus, the individual's definition of the event is more likely to be congruent with his or her attitude toward the object in cases involving strong object–evaluation associations than in cases involving weak ones. Consequently, behavior is more likely to follow from a definition of the event that is attitudinally based in the former cases than in the latter.

An additional implication of the present findings for the prediction of behavior from attitude is worth noting. The first two experiments identified attitudes that involved object–evaluation associations sufficiently strong to produce automatic activation. This was accomplished by assessing latency of response to an attitudinal inquiry. The findings imply that such a measure is a fairly good approximation of the likelihood of automatic activation of the evaluation upon mere observation of the attitude object. That is, how long it takes to respond to an attitudinal inquiry is reflective of the likelihood that the attitude will be activated spontaneously upon one's encountering the attitude object.

The very simplicity of the latency measure makes it attractive and potentially feasible for use in surveys in which one is concerned with the prediction of behavior from attitude. Predictive power should be enhanced by taking into account the accessibility of the respondents' attitudes. This notion was tested in a recent field investigation concerning attitudes toward Ronald Reagan (Fazio & Williams, 1985). Approximately 3½ months before the 1984 presidential election, attitudes toward Ronald Reagan and the accessibility of those attitudes (as measured by response latency to the attitudinal inquiry) were assessed. Attitudes were found to be far more predictive of later voting behavior among those respondents who were able to respond relatively quickly to the attitudinal inquiry than among those who responded relatively slowly.

Thus, there appear to be important implications of automatic attitudinal activation for both persuasion and attitude–behavior consistency. More generally, the present conceptualization and findings are relevant to the very functionality of attitudes. Attitude theorists have long considered one of the major functions served by attitudes to be that of organizing and structuring a rather chaotic universe of objects (Katz, 1960; Smith, Bruner, & White, 1956). An attitude is presumed to provide "a ready aid in 'sizing up' objects and events in the environment" (Smith et al., 1956, p. 41). The degree to which a given attitude actually fulfills this object appraisal function would appear to depend on the likelihood that the attitude is activated automatically when the individual observes the attitude object and, hence, on the strength of the object–evaluation association. Attitudes involving a strong association are highly functional. They free the individual from the processing required for reflective thought about his or her evaluation of the object and, through the process outlined earlier, can guide the individual's behavior in a fairly automatic manner. Thus, the individual is freed from much of the effort of having to engage in deliberate reasoning processes before behaving toward the object in question.

One final implication of the present research concerns attitude measurement. The priming procedure that was used to study the automatic activation of attitudes has the potential of serving as an unobtrusive measure of attitude. Attitude researchers have long searched for a means of assessing attitudes that was not subject to the respondent's self-presentational concerns with social desirability. Despite assurances of anonymity and confidentiality, individuals are not always completely honest when responding to an attitude survey. The most frequently used technique to deal with such assessment concerns has been the bogus pipeline (Jones & Sigall, 1971). This technique relies on an elaborate attempt to convince the respondents that a presumed apparatus has the capacity to discern their true attitudes from various physiological data that are ostensibly being monitored. Because individuals believe that it is not possible to deceive the machine, they respond more truthfully to the survey questions. In contrast, the present approach provides a means of obtaining an estimate of the attitude in a situation in which the individual need not be at all aware that his or her attitude is being assessed. Recall that during the critical priming phase of the experiments the subject was never asked to consider his or her attitude toward the object in question. Yet, it is possible to ascertain from the data the degree to which positive or negative evaluations are activated when the attitude object is encountered. Furthermore, given its automatic nature, such activation is presumably both inescapable and impossible for the individual to control. Whether the present technique is actually useful in measuring attitudes concerning objects or issues for which self-reports might be suspect remains to be tested. However, the present findings suggest that it may be a viable approach.

References

Abelson, R. P., Kinder, D. R., Peters, M. D., & Fiske, S. T. (1982). Affective and semantic components in political person perception. *Journal of Personality and Social Psychology, 42,* 619–630.

Anderson, J. R., & Pirolli, P. L. (1984). Spread of activation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* 791–798.

Converse, P. E. (1970). Attitudes and non-attitudes: Continuation of a dialogue. In E. R. Tufte (Ed.), *The quantitative analysis of social problems* (pp. 168–189). Reading, MA: Addison-Wesley.

deGroot, A. (1983). The range of automatic spreading activation in word priming. *Journal of Verbal Learning and Verbal Behavior, 22,* 417–436.

deGroot, A., Thomassen, A., & Hudson, P. (1982). Associative facilitation of word recognition as measured from a neutral prime. *Memory and Cognition, 10,* 358–370.

Fazio, R. H. (in press). How do attitudes guide behavior? In R. M. Sorrentino & E. T. Higgins (Eds.), *The handbook of motivation and cognition: Foundations of social behavior.* New York: Guilford Press.

Fazio, R. H., Chen, J., McDonel, E. C., & Sherman, S. J. (1982). Attitude accessibility, attitude–behavior consistency, and the strength of the object–evaluation association. *Journal of Experimental Social Psychology, 18,* 339–357.

Fazio, R. H., Powell, M. C., & Herr, P. M. (1983). Toward a process model of the attitude–behavior relation: Accessing one's attitude upon mere observation of the attitude object. *Journal of Personality and Social Psychology, 44,* 723–735.

Fazio, R. H., & Williams, C. (1985). *Attitude accessibility as a moderator of the attitude–perception and attitude–behavior relations: An investigation of the 1984 presidential election.* Unpublished manuscript, Indiana University.

Fazio, R. H., & Zanna, M. P. (1978a). Attitudinal qualities relating to the strength of the attitude–behavior relationship. *Journal of Experimental Social Psychology, 14,* 398–408.

Fazio, R. H., & Zanna, M. P. (1978b). On the predictive validity of attitudes: The roles of direct experience and confidence. *Journal of Personality, 46,* 228–243.

Fiske, S. T. (1982). Schema-triggered affect: Applications to social perception. In M. S. Clark & S. T. Fiske (Eds.), *Affect and cognition: The 17th Annual Carnegie Symposium on Cognition* (pp. 55–78). Hillsdale, NJ: Erlbaum.

Fiske, S. T. & Pavelchak, M. A. (in press). Category-based versus piecemeal-based affective responses: Developments in schema-triggered affect. In R. M. Sorrentino & E. T. Higgins (Eds.), *The handbook of motivation and cognition: Foundations of social behavior.* New York: Guilford Press.

Hovland, C. I. (1959). Reconciling conflicting results derived from experimental and survey studies of attitude change. *American Psychologist, 14,* 8–17.

Hovland, C. I., Harvey, O. J., & Sherif, M. (1957). Assimilation and contrast effects in reactions to communication and attitude change. *Journal of Abnormal and Social Psychology, 55,* 244–252.

Jones, E. E., & Sigall, H. (1971). The bogus pipeline: A new paradigm for measuring affect and attitude. *Psychological Bulletin, 76,* 349–364.

Jonides, J., & Mack, R. (1984). On the cost and benefit of cost and benefit. *Psychological Bulletin, 96,* 29–44.

Katz, D. (1960). The functional approach to the study of attitudes. *Public Opinion Quarterly, 24,* 163–204.

McKoon, G., & Ratcliff, R. (1980). Priming in item recognition: The organization of propositions in memory for text. *Journal of Verbal Learning and Verbal Behavior, 19,* 369–386.

McNamara, T. P., Ratcliff, R., & McKoon, G. (1984). The mental representation of knowledge acquired from maps. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* 723–732.

Neely, J. H. (1976). Semantic priming and retrieval from lexical memory: Evidence for facilitatory and inhibitory processes. *Memory and Cognition, 4,* 648–654.

Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General, 106,* 225–254.

Norman, R. (1975). Affective-cognitive consistency, attitudes, conformity, and behavior. *Journal of Personality and Social Psychology, 32,* 83–91.

Powell, M. C., & Fazio, R. H. (1984). Attitude accessibility as a function of repeated attitudinal expression. *Personality and Social Psychology Bulletin, 10,* 139–148.

Ratcliff, R., & McKoon, G. (1978). Priming in item recognition: Evidence for the propositional structure of sentences. *Journal of Verbal Learning and Verbal Behavior, 17,* 403–417.

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

Sample, J., & Warland, R. (1973). Attitude and prediction of behavior. *Social Forces, 51,* 292–304.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review, 84,* 1–66.

Sherif, C. W., Kelly, M., Rodgers, H. L., Sarup, G., & Tittler, B. I. (1973). Personal involvement, social judgment, and action. *Journal of Personality and Social Psychology, 27,* 311–328.

Sherif, M., & Cantril, H. (1947). *The psychology of ego involvements.* New York: Wiley.

Shiffrin, R. M. (in press). Attention. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, & R. D. Luce (Eds.), *Stevens' handbook of experimental psychology* (2nd ed.). New York: Wiley.

Shiffrin, R. M., & Dumais, S. T. (1981). The development of automatism. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 111–140). Hillsdale, NJ: Erlbaum.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review, 84,* 127–190.

Smith, M. B., Bruner, J. S., & White, R. W. (1956). *Opinions and personality.* New York: Wiley.

Wood, W. (1982). Retrieval of attitude-relevant information from memory: Effects on susceptibility to persuasion and on intrinsic motivation. *Journal of Personality and Social Psychology, 42,* 798–810.

Zanna, M. P., & Rempel, J. K. (1984). *Attitudes: A new look at an old concept.* Paper presented at the Conference on Social Psychology of Knowledge, Tel Aviv.

Received May 9, 1985
Revision received September 8, 1985 ■