The Generality of the Automatic Attitude Activation Effect

John A. Bargh and Shelly Chaiken
New York University

Rajen Govender
University of the Western Cape, Bellville, South Africa

Felicia Pratto
Stanford University

Fazio, Sanbonmatsu, Powell, and Kardes (1986) demonstrated that Ss were able to evaluate adjectives more quickly when these adjectives were immediately preceded (primed) by attitude objects of similar valence, compared with when these adjectives were primed by attitude objects of opposite valence. Moreover, this effect obtained primarily for attitude objects toward which Ss were presumed to hold highly accessible attitudes, as indexed by evaluation latency. The present research explored the generality of these findings across attitude objects and across procedural variations. The results of 3 experiments indicated that the automatic activation effect is a pervasive and relatively unconditional phenomenon. It appears that most evaluations stored in memory, for social and nonsocial objects alike, become active automatically on the mere presence or mention of the object in the environment.

The research we report in this article concerns the extent to which attitudes may be activated automatically on mere observation of the attitude objects to which they correspond. This phenomenon is central to Fazio's (1986, 1989, 1990) *attitude accessibility model,* a theory of the process by which attitudes guide behavior. Consistent with the attitude accessibility model, we demonstrate that attitudes are capable of automatic activation. Yet, the previous experimental tests of the automatic activation effect have obtained it primarily for a given subject's most highly accessible or strongest attitudes as opposed to the subject's least accessible or weakest attitudes (Fazio, Sanbonmatsu, Powell, & Kardes, 1986). We present data indicating that such automatic activation occurs for most of a subject's attitudes.

According to the attitude accessibility model, the first step in the process by which attitudes guide behavior is *attitude activation,* the retrieval of one's evaluation of the attitude object from memory. Once activated, the attitude influences perception of the attitude object and the situation in which it is encountered, and these perceptions, in turn, influence subsequent behavior toward the attitude object. The activation step of this model is critical, because only activated attitudes can be expected to guide subsequent information processing and behavior (Fazio, 1986, 1989, 1990).

What, then, determines the likelihood that a person's attitude will be activated on encountering the attitude object? According to the model, the prime determinant of activation is *associative strength*—the strength of the association (in memory) between an attitude object (e.g., *television*) and an evaluation (e.g., *bad*). This is because associative strength is postulated to determine an attitude's *chronic accessibility,* the likelihood that it will be activated on the person's exposure to the attitude object (or to other attitude-relevant cues; see Fazio, 1989).

Although not antithetical to models that view the attitude–behavior relation as the product of controlled, deliberative reasoning processes (e.g., Ajzen & Fishbein, 1980; see Fazio, 1990), the attitude accessibility model asserts that attitudes can affect behavior without any effort or intention on the part of the indi-

Editor's Note. Two of the authors are Associate Editors of this journal; therefore, this article was subjected to a special review process to avoid any potential conflict of interest. I asked Anthony Greenwald, a former Editor of *JPSP,* to handle the article. He solicited the reviews and made the recommendations that the article be published. I am grateful to him for taking on this chore.

This research was supported by National Institute of Mental Health Grants R01-MH43265 to John A. Bargh and R01-MH43299 to Shelly Chaiken. Preparation of the article was facilitated by a fellowship from the Alexander von Humboldt Foundation to John A. Bargh. We thank Bess Buckley, Marc Kossmann, Gloria Lilliard, and Paula Raymond for their assistance with the data collection and analysis and Icek Ajzen, Niall Bolger, Alice Eagly, Roger Giner-Sorolla, Tory Higgins, Akiva Liberman, Gordon Logan, Diane Ruble, Lisa Spielman, and several anonymous reviewers for their comments on earlier versions of this article. We also thank Jack Cohen for his invaluable advice regarding the regression analysis and Russell Fazio and David Sanbonmatsu for supplying us with their prime frequency data and stimulus materials.

Correspondence concerning this article should be addressed either to John A. Bargh or to Shelly Chaiken, Department of Psychology, New York University, Seventh Floor, New York, New York 10003.

vidual, that is, in a spontaneous, automatic manner (see Eagly & Chaiken, in press, for discussion). Indeed, it is this unique feature of the model that has been emphasized in research (for reviews, see Fazio, 1986, 1989, 1990).

While promoting the view of an automatic attitude-to-behavior sequence, theoretical statements of the attitude accessibility model have cautioned that not all attitudes should be capable of influencing behavior in this manner. Only those attitudes that are capable of being automatically activated on mere observation of the attitude are hypothesized to possess this power, and the likelihood of such activation, according to the model, depends on associative strength (Fazio, 1986, 1989, 1990). The attitude accessibility model conceptualizes associative strength as a continuum, with those attitudes of relatively high associative strength capable of achieving automaticity:

> At the lower end of the continuum is the nonattitude. No a priori evaluation of the attitude object exists in memory. As we move along the continuum, an evaluation does exist and the strength of the association between that evaluation and the object and, hence, the chronic accessibility of the attitude, increases. In the case of a weak association, the attitude can be retrieved via an effortful, reflective process but is not capable of automatic activation. At the upper end of the continuum is a well-learned, strong association that is likely to be activated automatically upon mere observation or mention of the attitude object. (Fazio, 1989, pp. 159–160)
>
> Given that automatic processes require such well-learned responses, it appears doubtful that automatic activation is likely for all attitudes that an individual might hold. Only for well-learned ones is the expectation of automatic activation even a possibility. (Fazio et al., 1986, p. 229)

The automatic activation hypothesis was tested in three experiments by Fazio et al. (1986) that used a sequential priming procedure to assess automaticity (e.g., Neely, 1977). Because we used the same basic procedure in the present studies, it is necessary to describe it in some detail here. On each trial, the name of an attitude object (e.g., *landlords*) was briefly presented (for 200 ms) on a computer-driven display. After a 100-ms blank-screen interval, an adjective appeared in the same screen location, the subjective effect of which was to overwrite the attitude object stimulus. The subject's task was to indicate whether the displayed adjective was "good" or "bad" in meaning, and the latencies of these judgments were recorded. The logic of the design was that to the extent that presentation of the attitude object name activated the evaluation associated with the attitude object, this evaluation (good or bad) would then influence how quickly subjects could correctly classify the target adjective as positive or negative in meaning. If the adjective was of the same valence as the attitude object prime, responses should have been faster (i.e., facilitated) relative to a baseline (nonword letter string) prime condition. Conversely, if the adjective and prime were of opposite valence, responses should have been relatively slower.

The stimulus onset asynchrony (SOA) of 300 ms between the attitude object prime and target adjective presentations is a critical feature of the paradigm (one that we did not vary in the reported research). As Fazio et al. (1986) argued, it is too brief an interval to permit subjects to develop an active expectancy or response strategy regarding the target adjective that follows; such conscious and flexible expectancies require at least 500 ms to develop and to influence responses in priming tasks (Neely, 1977; Posner & Snyder, 1975). Given an SOA of 300 ms, then, if presentation of an attitude object prime influences response time to a target adjective, it can only be attributed to an automatic, unintentional activation of the corresponding attitude. Importantly, Fazio et al. (1986) included a 1000-ms SOA condition in their Experiments 2 and 3 and, consistent with their expectation that subjects in this condition would have sufficient time to actively disregard the influence of the prime, showed that the pattern of facilitation and inhibition predicted from the automatic activation hypothesis did not occur (see also Neely, 1977).

The Fazio et al. (1986) attitude object prime stimuli were chosen idiographically in two of the studies and normatively in the third. Before the priming task phase of the first 2 experiments, subjects indicated their attitudes toward each of a large set of attitude objects (e.g., 92 in Experiment 2) by pressing a *good* key or a *bad* key. The latency and valence of these evaluations were used to select 16 attitude objects for each subject. These objects represented four types of primes (Strong–Weak × Good–Bad) and corresponded to each subject's most quickly expressed (i.e., "strongest") or most slowly expressed (i.e., "weakest") positive and negative attitude judgments. Thus, the strength of a person's attitude was conceptualized in terms of associative strength and accessibility and operationalized in terms of evaluation latency. In Experiment 3, relatively weak attitude objects (8 good and 8 bad) were preselected for all subjects from those relatively slowly evaluated in Experiments 1 and 2 but for which subjects showed high consensus in their evaluations. Just before the priming task in this experiment, associative strength was manipulated by having subjects either evaluate each attitude object five times (creating temporarily accessible attitudes) or make a grammatical judgment about them five times (leaving the attitude relatively weak).

The results of Experiments 1 and 2 were virtually identical. Whereas a significant automatic activation effect was observed when subjects were primed with attitude objects toward which they presumably possessed strong, highly accessible attitudes, this effect did not prove reliable when the prime stimuli were attitude objects toward which subjects presumably possessed weak, inaccessible attitudes. Experiment 3 of this series also obtained a reliably greater automatic activation effect for attitude object primes toward which subjects' attitudes had been temporarily strengthened by the repeated expression manipulation. However, in contrast with the null findings observed in the weak attitude conditions of the first 2 studies, the automatic activation effect did prove reliable for Experiment 3's control attitude object stimuli, which had been among the slowest evaluated in the first 2 experiments and toward which subjects were presumed to hold weak, relatively inaccessible attitudes.

How General Is the Automatic Activation Effect?

Fazio et al. (1986) concluded that their results supported the hypothesis that the likelihood of automatic attitude activation depends on the associative strength of one's attitude. However, there are a few causes for concern about this interpretation. First, the Experiment 3 control condition results would seem inconsistent with the idea that only highly accessible attitudes can achieve automaticity (Fazio, 1986, 1989, 1990). Moreover, Sanbonmatsu, Osborne, and Fazio (1986) also obtained a reliable automatic activation effect for relatively weak attitude object primes, although the effect was again reliably greater for strong attitude object primes.

Second, because the Fazio et al. (1986; Sanbonmatsu et al., 1986) experiments examined only each subject's four strongest (fastest evaluated) and four weakest (slowest evaluated) good and bad attitudes, it is unclear how general the automatic activation effect is across attitude objects of variable associative strengths. In harmony with the accessibility model, perhaps the effect is restricted to subjects' strongest attitudes. On the other hand, these studies' designs do not rule out the possibility that the automatic activation effect holds for all but the subject's weakest attitudes. Importantly, if all attitudes except the very weakest ones (i.e., nonattitudes) showed the automaticity effect, associative strength might not be the determining factor—the existence of any object–evaluation association might be all that is necessary.

Conditional Automaticity

The importance of the automatic attitude activation effect for the attitude accessibility model is that it apparently demonstrates that some attitudes (i.e., the strongest) become active "upon mere observation or mention of the attitude object" (Fazio, 1989, p. 160). Indeed, this was the explicit goal of the Fazio et al. (1986) experiments. Yet, as described below, in actuality the subject's task in these experiments required more processing of the attitude object primes than just minimal attention to them. Many if not most automatic phenomena in fact depend on certain preconditions to occur; for example, the automatic activation of a trait construct by a relevant behavioral event often does not occur unless that construct has been recently activated (primed) by the situational context, and automated skills (from typing to forming an impression of someone) are not engaged in without the intention or goal to do so (see Bargh, 1989, 1992). Thus, more than the mere observation of the object may be necessary for the effect to be obtained.

Temporary Versus Chronic Accessibility

The attitude assessment task in the Fazio et al. (1986) experiments required subjects to intentionally and consciously evaluate attitude objects. Because this task was performed immediately before the priming task that tested for automaticity, the effects demonstrated may have depended on subjects having just thought about their attitudes toward the prime stimuli. In other words, the automaticity apparently demonstrated might, instead, be due to the priming of object–evaluation associations by the attitude assessment task; such temporary activation effects caused by priming have been found to mimic automatic, chronic effects (Bargh, Bond, Lombardi, & Tota, 1986; Bargh, Lombardi, & Higgins, 1988; Spielman & Bargh, 1991). If demonstration of the automatic attitude effect depends on the subject having just consciously and intentionally evaluated the attitude object, the effect may well require more than just the mere presence of the attitude object in the environment.

"Memory Word" Instructions

In the Fazio et al. (1986) priming task, subjects were instructed to hold the attitude object prime in memory as a "memory word" throughout the trial and, after giving their evaluation of the target adjective on that trial, to repeat the prime word aloud. Fazio et al. (1986) included this procedure to ensure that subjects paid attention to the prime word on each trial. However, it is not a standard procedure in priming studies testing for the presence of a stimulus-driven automatic process (e.g., Logan, 1980; Neely, 1977; Posner & Snyder, 1975), because such a process must be shown not to require any intentional, conscious processing to occur (see Logan, 1989). By rehearsing the prime stimuli during the priming trials, subjects gave these stimuli more conscious attention and processing than that involved in merely noticing their presence. The automatic activation effect found by Fazio et al. therefore also may depend on the paradigm's memory word instructions. If so, it might not generalize to settings in which people do not concentrate so intensively on the attitude object when they encounter it.

Present Experiments

We conducted three experiments to determine the generality of the automatic activation effect, by exploring the range of the effect across the set of 92 attitude objects used in the Fazio et al. (1986) research and by testing for the presence of the effect when alterations to the experimental paradigm are made. As a necessary first step in assessing the range of the effect across the 92 attitude object stimuli, we collected normative data on those characteristics of the stimuli that potentially were related to their ability to become activated automatically.

Collection of Normative Data

There were three phases of normative data collection: a questionnaire administered at a mass testing session, a separate questionnaire administered to a smaller sample, and an experimental test of the stability of attitude evaluations over time. For the sake of clarity, the three methods are described together, followed by a description of the results we obtained.

Method

Evaluation latencies and extremity ratings. What features of attitude object stimuli might be related to the automatic activation effect? We first assessed the extremity of subjects' attitudes toward each of the 92 attitude objects, an indicator of attitude strength that may be correlated with but is not redundant with attitude accessibility. Indeed, previous research shows that greater attitudinal extremity is associated with faster evaluation latencies (e.g., Fazio & Williams, 1986; Judd & Kulik, 1980; Powell & Fazio, 1984; Wyer & Gordon, 1984). For example, Powell and Fazio (1984) reported a correlation of .30 between these variables, and Fazio and Williams (1986) reported an average correlation of .53. In fact, the magnitude of this relation led Fazio and Williams to control for extremity in their investigation of the impact of accessibility on the predictability of behavior from attitudes.

As part of a mass testing session at the beginning of the semester, 274 introductory psychology students at New York University (NYU) completed a Semantic Evaluations Questionnaire (SEQ). This questionnaire contained 141 items: the 92 attitude object stimuli used in the

Fazio et al. (1986) Experiment 2 and 49 trait terms, selected on the basis of Anderson's (1968) norms to represent the range of likability, that served as filler items. Subjects were instructed to give their personal evaluation of each of the "objects, concepts, and qualities" listed by circling one number on an 11-point scale anchored by *extremely bad* (−5) and *extremely good* (5). Two random orders of the 141 items were developed, and subjects received one or the other on a random basis.

Evaluation latencies and consistency of evaluation over time. It is also possible that the consistency with which people evaluate attitude objects is related to the automatic activation effect. Like extremity, consistency may covary with associative strength or accessibility but may nonetheless represent an independent indicator of attitude strength (Doll & Ajzen, 1991). The Fazio et al. (1986) Experiment 3 used attitude object primes that were preselected to have "the longest average response latencies" among the Experiment 1 and 2 primes that had been "endorsed with near unanimity across the subjects" (p. 234). The low attitude strength control condition of this study yielded a reliable automatic activation effect for these prime stimuli.

Why should attitude objects be capable of automatic activation by virtue of high consensus regarding their goodness or badness? It is likely that attitude objects for which there is evaluative consensus across subjects are also those that the individual subject evaluates consistently over time. For example, McCloskey and Glucksberg (1978) found that the amount of agreement across subjects in classifying stimuli in terms of categories (e.g., lamp = FURNITURE?) was highly related to the degree of consistency over time shown by individual subjects. Consistency of evaluation over time within the individual should be important to the development of automatic attitudes, because the automatization of mental pathways requires both considerable frequency of pathway activation and consistency of activating that particular path rather than another (Fisk & Schneider, 1984; Milner, 1957; Posner, 1978; Schneider & Fisk, 1982; Schneider & Shiffrin, 1977). For example, Schneider and Fisk (1982) found automatic process development in a visual search task to be a multiplicative function of practice and consistency: Unless a given stimulus appeared with a moderately high probability of being the target (as opposed to a distracter) of the search, even thousands of trials of practice produced no trace of an automatic detection capability.

To examine the potential role of consistency of evaluation in producing the automatic attitude activation effect, we conducted a study designed to assess the relative consistency of subjects' attitudes toward the 92 attitude stimuli used in the Fazio et al. (1986) experiments. Three to 6 weeks after the mass testing session in which extremity ratings were obtained, each of a sample of 30 of these students participated in two laboratory sessions, 2 days apart, in partial fulfillment of a course requirement. Subjects sat in front of a cathode-ray tube (CRT) display, which was under program control of an Apple II Plus microcomputer. On the table in front of the subject was a two-button response box that served as an input device to the computer. Subjects rested their hands comfortably on the box so that the index finger of their left hand (right hand) was poised near the button on the left side (right side) of the box. The left button was labeled *bad* and the right button, *good.*

Subjects participated one at a time. At the first session, the subject was shown into the 2.7-m × 3-m experimental room and given both verbal (by the experimenter) and visual (by CRT display) instructions. Subjects learned that a series of words would be presented on the screen and that their task was to indicate their evaluation of the object represented by each word by pressing either the *good* or the *bad* button. They were told to respond quickly but accurately. Before the experimental trials, subjects completed 10 practice trials that presented attitude objects similar in valence and content to those used on the experimental trials.

After the practice trials, the experimenter left the room and pressed a key on the computer that caused the experimental trials to begin. One of two random orderings of the 92 stimuli was presented to the subject. Each of the attitude objects was presented in the middle of the computer screen and remained on the screen until the subject pressed one of the response buttons. The computer recorded the response (good or bad) and its latency (in milliseconds) for each trial. As in the Fazio et al. (1986) procedure, there was a 3-s pause before the next trial. After completing the experimental trials, subjects were reminded about their appointments for the second session and excused.

At the second session subjects again evaluated the 92 attitude objects. The procedure was the same as the one used 2 days earlier, except that the random ordering of stimuli not used in Session 1 was used in Session 2 (order of stimuli was counterbalanced across subjects and had no influence on any of the data reported later). After completing the evaluation task, subjects responded to whichever form of the SEQ they had not completed during the mass testing session earlier in the semester. Afterward, they were thanked for participating, debriefed, and excused.

Evaluation latencies and attitude ambivalence. Because both frequency and consistency of attitude object evaluation should contribute to the strength of the object–evaluation association and, thus, the accessibility of the evaluation, two different types of attitude objects may not show the automatic activation effect—those with which the subject has had little or no direct experience (low frequency of evaluation) and those with which the subject has had considerable experience but toward which he or she has ambivalent feelings (low consistency of evaluation). The former case is that of the nonattitude (Converse, 1970), which is not really there in memory to become active in the first place and so should not be capable of automatic activation under any circumstances (Fazio, 1989).

Ambivalent attitudes, on the other hand, correspond to attitude objects that have rather strong links in memory to both good and bad evaluations (Kaplan, 1972). Thus, on presentation of the attitude object, both the good and the bad evaluations presumably would become active. Yet it may take subjects longer to give their attitude in the assessment task because of response competition: They must choose one or the other response (good vs. bad), hence one of the activated responses must be made and the other inhibited (e.g., Logan, 1980; Shallice, 1972). This act of suppressing a competing response requires attentional resources and hence time as well (Garner, 1962; Katz, 1981, pp. 361–372). The extent to which competing response tendencies slow response times no doubt depends on their relative strengths. Polarization, which is inversely related to ambivalence, reflects this property (see following discussion).

Accordingly, we asked an additional group of 32 introductory psychology students to complete a questionnaire measuring the ambivalence and polarization of their attitudes toward each of the 92 attitude objects and also to complete the attitude assessment task, following the same procedure as for the consistency study. The questionnaire was based on Kaplan's (1972) method of including separate unipolar positive and negative scales for each item to be rated. Respondents are first instructed that people often have both positive and negative feelings toward issues, people, and objects, and so for each item on the questionnaire, they are to rate the extent to which they have positive feelings toward the item and then, separately, the extent to which they have negative feelings. Both the positive and the negative scale consisted of four possible responses: *not at all, slightly, quite,* and *extremely.* From these separate scales indexes of ambivalence and polarization of attitudes can be computed (see following discussion). Two orderings of the 92 attitude objects were created; also, either the positive scale or the negative scale came first for all items, and the subject completed either the ambivalence questionnaire or the attitude-assessment task first. All of these factors were counterbalanced in the design.

Results

Computation of norms. We first computed, for each of the 92 attitude objects, the mean across subjects of (a) evaluation ratings on the SEQ at the mass testing ($N = 274$), (b) evaluation latency at the first laboratory session ($n = 30$), (c) evaluation latency at the second laboratory session ($n = 30$), (d) evaluation ratings at the end of the second laboratory session ($n = 30$), (e) ambivalence ratings by the separate sample of 32 subjects, (f) polarization ratings by those subjects, and (g) their evaluation latencies.

We computed the ambivalence and polarization indexes for each of the 92 attitude objects, following Kaplan's (1972) procedures. First, responses on both the positive and the negative feelings scales were coded as 0 (*not at all*), 1 (*slightly*), 2 (*quite*), or 3 (*extremely*). Next, the subject's degree of ambivalence toward a given attitude object was computed by taking the sum of the positive and negative ratings of the attitude object (this represents the total amount of affect toward the object, regardless of valence) and then subtracting the absolute value of the difference between the two scales. Thus, the ambivalence index represents the amount of precisely counterbalancing negative and positive affect toward the object. It can range from 0, in the case in which the response to one or both scales was *not at all,* to 6, in the case in which the subject had both extreme positive and extreme negative feelings toward the object.

The polarization index was calculated by taking the absolute value of the difference between the positive and the negative scales; thus, it reflects the extent to which the subject has stronger positive (negative) than negative (positive) feelings toward the object. Polarization scores can range from 0 (when the positive and negative ratings are equal in extremity) to 3 (when one rating is *extremely* and the other is *not at all*).

We next computed for each attitude object the percentage of subjects (out of 30) who gave the same evaluation of it during the second speeded evaluation task as during the first (i.e., good–good or bad–bad). This served as our operationalization of the consistency with which attitudes toward that object are held, based on our assumption that attitude objects with high consensus of evaluation across subjects are evaluated consistently over time by an individual subject as well (see earlier discussion). The mean number of such changes in evaluation per attitude object was 2.6 (corresponding to a mean of 91.3% consistency), with a standard deviation of 2.3 (7.7%), median of 2.0 (93.3%), and range of 0–8 (100% to 73.3%).

Finally, because the frequency of a word in the language as well as its length can affect response latencies (see Whaley, 1978), we also recorded this information for each of the 92 attitude objects.

Stability of indexes. Mean evaluation latencies proved highly stable across the two testing sessions of the consistency study, with $r = .80$ for the raw latencies across the 92 attitude objects and a rank-order correlation of .79 (both $ps < .01$). To eliminate the influence of retesting, we calculated the mean evaluation latencies for the 92 attitude objects by using data from only the first of the consistency study sessions ($n = 30$) combined with data from the ambivalence study ($n = 32$). The mean latencies for the 92 stimuli from the first consistency attitude assessment and the ambivalence study assessment task were highly correlated as well, $r = .62$, $p < .001$.

Evaluation ratings of the 92 attitude objects (on the −5 to 5 scale) were highly stable, with $r = .99$ ($p < .001$) between the mean ratings by the 30 consistency experiment subjects, taken at the mass testing and then 3–6 weeks later at the end of the second session. Moreover, both sets of rating means correlated highly (both $rs = .97$, $ps < .001$), with the means based on the entire mass testing sample of 274 students. Subsequent analyses used the means for this larger sample as they provide the more stable estimates of the valence (i.e., dichotomous positive vs. negative evaluation) and extremity of evaluation (i.e., absolute value of the mean rating) of the 92 attitude objects.

Correlational analyses. The word length, word frequency, mean evaluation latency, mean evaluation, mean consistency of evaluation, mean ambivalence, and mean polarization scores of each of the 92 attitude objects appear in the Appendix. Table 1 displays the intercorrelations among these factors (with valence considered separately from extremity of evaluation), computed across the 92 attitude objects.[1]

All of the predictor variables were significantly correlated with mean evaluation latency. Attitude objects represented by longer words were responded to more slowly. Also, the more frequent the word in the language, the more quickly subjects responded to it (i.e., shorter latencies). The more ambivalent an attitude toward an object, the longer subjects took in giving their evaluation of it. Attitude polarization was also reliably related to evaluation latency, with subjects responding faster to an attitude object the more polarized their attitude toward it. Finally, positive evaluations were made more quickly than negative evaluations, but this may have been due merely to a greater readiness to respond *good* rather than *bad.*[2]

Attitude extremity and consistency (consensus) were the features most highly related to evaluation latency. As in previous research, the more extremely held an attitude, the faster its corresponding attitude object was evaluated. Moreover, as predicted, the consistency with which an attitude object was evaluated was strongly related to evaluation latency—the more consistent the attitude, the faster the evaluation latency.

In addition to providing normative data on the set of 92 attitude objects, these studies have also established that several factors other than associative strength (which determines accessibility) may influence the speed with which people give their evaluations of attitude objects. These include the length and frequency in the language of the word that represents the atti-

[1] Note that Table 1 and the Appendix present an additional consistency index, labeled *consensus,* which is based on data from all experiments reported in this article and is described in the Overall Regression Analysis section. In addition, the latency means used in the correlational analysis and presented in the Appendix are based on the attitude assessment data from Experiments 1–3 as well as the normative consistency and ambivalence experiments.

[2] The high negative correlation between ambivalence and polarization shown in Table 1 is to be expected given that a high ambivalence score necessarily means a low polarization score, and a high polarization score puts a low ceiling on the possible ambivalence score (see computation of norms discussion).

Table 1
Correlations Among Attitude Object Characteristics and Evaluation Latency, From Normative and Experimental Data (see Appendix)

Attitude object characteristic	1	2	3	4	5	6	7	8	Latency
1. Length	—	−.34**	.07	−.03	−.00	−.07	−.05	−.04	.25*
2. Frequency		—	.16	.06	.07	.28**	.13	.23*	−.24*
3. Ambivalence			—	−.90**	.34**	−.54**	−.27**	−.44**	.24*
4. Polarization				—	−.32**	.78**	.46**	.66**	−.43**
5. Valence					—	−.01	.21	.17*	−.38**
6. Extremity						—	.70**	.93**	−.69**
7. Consistency							—	.72**	−.74**
8. Consensus								—	−.72**

Note. *Consistency* refers to the percentage of subjects in the normative consistency study who gave the same evaluation of the attitude object at both experimental sessions. Consensus is a factor score that is based on data from subjects in all experiments; see Appendix, footnote e.
* $p < .025$. ** $p < .01$.

tude object. It is also possible that attitude consistency, extremity, ambivalence, and polarity are independent predictors of evaluation latency (see Eagly & Chaiken, in press). In terms of the attitude accessibility model, however, the relation of the extremity and consistency variables to evaluation response latency can be viewed as due to their effect on the strength of the object–evaluation association (but see Doll & Ajzen, 1992; Downing, Judd, & Brauer, in press).[3] Because attitude consistency theoretically should be a determinant of the automatic activation effect, and as it proved to be one of the strongest predictors of evaluation latency, we emphasized the role of attitude consistency in the first experiment.

Experiment 1

As noted earlier, the finding that relatively slowly but consistently evaluated attitude objects reliably produced the automatic activation effect in the Fazio et al. (1986) Experiment 3 had suggested to us that the effect might hold for a greater range of associative strengths than just the very fastest evaluated. That is, even relatively weak attitudes (as operationalized in terms of evaluation latency) appear to show the effect as long as they are consistently evaluated over time.

To test this hypothesis, we conducted a replication and extension of the Fazio et al. (1986) Experiment 2, with the inclusion of an additional set of "consistent" attitude object primes selected on the basis of our normative data (see Appendix). As shown in Table 2, these preselected sets of objects that our normative study subjects had rated with high consistency as either positive or negative in valence were characterized by mean evaluation latencies that fell across the entire middle range of the latency distribution. So as not to confound consistency with extremity, we selected only those attitude objects marked by relatively moderate attitudes; of the 16 attitude objects shown in Table 2, 7 were below the median of the extremity distribution, and all but 1 fell in the 2nd or 3rd quartiles of the distribution.

We assumed, on the basis of the substantial correlation between consistency and latency obtained in our normative research (see Table 1), that the subject's fast attitude object primes would correspond to highly consistently evaluated attitude objects and his or her slow attitude objects, to relatively inconsistently evaluated attitude objects. Therefore, under the hypothesis that consistently evaluated attitude objects would show the automatic activation effect, whether or not they are the most quickly evaluated by the subject, we predicted that the fast and the consistent attitude objects would show the effect, whereas the slow attitude objects would not. This predicted pattern would replicate the Fazio et al. (1986) findings but with the additional demonstration that the effect holds across the range of attitude objects that are relatively consistently evaluated by subjects.[4]

Method

Subjects. Twenty-three introductory psychology students at NYU participated in the experiment. Because the study was conducted the semester after the semester in which the normative data collection had occurred, none of these students had been previously asked for their attitudes toward the study's 92 attitude objects.

Materials and apparatus. The experimental room and computer/response box apparatus were identical to those described for the normative consistency experiment described earlier. For the initial attitude assessment task, four random orders of the 92 attitude objects were created, and each subject was randomly assigned to one order (this factor had no influence on the reported results and is not discussed further).

Twenty-eight adjectives served as the target stimuli in the priming task. Fazio et al.'s (1986) 10 positive adjectives (e.g., *outstanding*) and 10 negative adjectives (e.g., *horrible*) were supplemented by another 4 posi-

[3] Ambivalence and polarization, on the other hand, are difficult to relate to the strength of a single object–evaluation association in memory, because they refer to the relative strengths of separate positive and negative evaluations of the object.

[4] Because the Fazio, Sanbonmatsu, Powell, and Kardes (1986) Experiment 2 (and 3) demonstrated an automatic activation effect when the prime–target asynchrony (SOA) was 300 ms, but not when this interval was set at 1000 ms, only the former condition was represented in our experimental design (see introduction).

Table 2
Consistent/Slow Attitude Objects Used as Priming Stimuli for All Subjects in Experiments 1 and 2

Attitude object prime	Latency M (ms)	Latency %[a]	Consistency (%)[b]
Positively valenced			
Magazine	674	50	97
Clothes	683	57	97
Cake	641	30	100
Stereo	692	63	97
Strawberries	708	76	93
Snow	719	80	93
Gold	658	35	93
Aquarium	721	83	97
Negatively valenced			
Alcoholism	795	80	100
Hornet	743	50	93
Hangover	782	70	93
Litter	773	61	93
Cavities	745	52	93
Bombs	765	59	93
Guns	652	11	97
Garbage	742	46	97

Note. Attitude objects are listed in order of their priority of use in the experiments (see text).
[a] Percentages refer to how many of the 92 latency means were lower (faster). [b] Consistency scores are the percentage of subjects in the normative consistency study (n = 30) who responded with the same valence in their evaluation of the object at both experimental sessions.

tive and 4 negative adjectives that we generated (*beautiful, excellent, magnificent,* and *marvelous; miserable, hideous, dreadful,* and *painful*). Each of these targets was paired in the priming task with one of each of the seven types of priming stimuli (fast-good, fast-bad, slow-good, slow-bad, consistent-good, consistent-bad, and baseline), creating a total of 196 experimental trials. The baseline primes were the same three-letter strings (e.g., BBB) used by Fazio et al.

A single random ordering of the 196 trials was created, with the restriction that each target adjective appear once in each successive block of 28 trials. Each of the 28 primes (4 each of the 7 prime types) appeared in a different random order in each successive trial block, with the restriction that 2 of each prime type be paired with positive targets and 2 be paired with negative targets in each block. A final restriction was that each target adjective be paired with 1 prime from each of the 7 prime types over the course of the experimental trials. In this way, every target adjective appeared the same number of times in each of the seven priming conditions.

Procedure. The procedure was identical to that used in the Fazio et al. (1986) Experiment 2 (300-ms SOA conditions) with the exception of the addition of the consistent good and bad prime types. Subjects were shown into the experimental room and seated in front of the CRT display and response box. They were told that the experiment concerned judgments about words and that they would perform two different tasks. The first of these, constituting the attitude assessment phase, was explained to them, and some practice with the task was given with the experimenter present. The procedure for this task was the same as in our normative study. Subjects gave their evaluation of each of the 92 attitude objects by pressing either the *good* or the *bad* button as quickly as possible after seeing the name of the attitude object appear on the screen.

After the assessment task, subjects were given a 3-min break while the computer program sorted their *good* and *bad* responses in ascending order of latency. The program then selected the attitude objects to which the four fastest and four slowest *good* responses were made and to which the four fastest and four slowest *bad* responses were made. It then checked the predetermined sets of consistent *good* and consistent *bad* attitude objects (a) to ensure that the subject had given the same evaluation of them as we had presumed on the basis of our norms (all of the subjects had) and (b) for possible overlap with the objects selected as the fastest or slowest within each valence type. To ensure a strict replication of the Fazio et al. (1986) procedure, in case of such overlap the attitude object in question served as the fast or slow prime, and a replacement consistent prime of the same valence was selected by the program from the set shown in Table 2. The replacement was the consistent prime with the highest priority from the list that did not already appear as a fast or slow prime for that subject. The experiment's preselected consistent primes are shown in Table 2 in the priority order in which the computer program selected them. These attitude object primes were entered into the list of prime stimuli by the computer program in the manner described in the Materials and Apparatus section.

For the second experimental task, subjects were again instructed to evaluate (as quickly as possible) words presented to them on the CRT screen. However, as in the Fazio et al. (1986) procedure, subjects were told that their task would be more difficult this time, as they would also have to remember a memory word while making each judgment, and to recite that word aloud after making their evaluative judgment of the second word presented on each trial. A microphone with a cord leading out of the room was placed in front and to the side of the subject to enhance the belief that the experimenter would be monitoring the subject's recitations of the memory words. Ten practice trials with this more complex procedure were given before the experimental trials. After answering any questions, the experimenter left the room and started the computer program that presented the experimental priming trials.

Again as in the Fazio et al. (1986) procedure, the prime word was presented in the middle of the CRT screen for 200 ms, with a 100-ms blank-screen interval before the target adjective appeared. The target remained on the screen until the subject pressed one of the two response buttons, and there was a 4-s pause before the start of the next trial. At the conclusion of the 196 trials, the subject was debriefed and thanked for participating.

Results

Subjects made few errors in their evaluation of the target adjectives (mean rate = 2.2%); latencies corresponding to these responses were excluded from all analyses (see Fazio et al., 1986). To reduce the influence of outlier response latencies on the analyses, those over 2,500 (0.7%) were set equal to 2,500.

For each subject we first computed the mean response latency for each of the 14 cells of the design (the seven prime types crossed with target adjective valence). Next, we computed facilitation scores by subtracting the means in the positive target conditions from the baseline prime–positive target mean and by subtracting the negative-target condition means from the mean for the baseline prime–negative target condition. As in the Fazio et al. (1986) studies, therefore, these facilitation scores represent the increase in evaluation latency for the identical set of positive or negative adjectives caused by the presentation of the given prime type immediately before those adjectives, relative to the condition in which meaningless letter strings were presented.

The facilitation scores for the six experimental conditions are shown in Figure 1. The 3 (prime type: fast vs. slow vs. consistent) × 2 (prime valence: positive vs. negative) × 2 (target valence: positive vs. negative) analysis of variance (ANOVA) on these scores yielded a reliable two-way interaction between prime valence and target valence, $F(1, 22) = 32.49$, $p < .001$, $MS_e = 8199.77$, accounting for 4.3% of the total variance, as well as a reliable three-way interaction, $F(2, 44) = 3.28$, $p < .05$, $MS_e = 10{,}021.18$, which accounted for 1.1% of the total variance. Given the reliable three-way interaction, we followed the Fazio et al. (1986) analysis strategy by next testing for the simple interaction between prime valence and target valence within each of the three prime types. As the pattern of means predicted by the automatic attitude activation effect is that of relative facilitation of latencies when the prime valence matches the target valence, and relative inhibition of latencies when prime and target valence mismatch, the presence of this pattern along with a reliable Prime Valence × Target Valence interaction indicates the presence of the effect.

As Fazio et al. did (1986, Experiments 1 and 2, 300-ms SOA condition), we obtained a reliable Prime Valence × Target Valence interaction for the fast prime condition, $F(1, 44) = 7.24$, $p = .01$, but not for the slow prime condition, $F(1, 44) = 1.70$, $p = .15$. Most important, the simple two-way interaction was also present in the consistent prime condition, $F(1, 44) = 23.98$, $p < .001$.

Discussion

In this experiment we replicated the Fazio et al. (1986) demonstration of an automatic attitude activation effect for attitudes that can be expressed very quickly and also the lack of the effect for attitudes that require a relatively long time to be expressed. In addition, however, we demonstrated that this automatic activation effect occurs for attitudes that were not relatively quickly expressed by subjects in the normative study, but which nonetheless had been consistently expressed by them. These results therefore confirm our hypothesis that consistently expressed attitudes would show the automatic activation effect even when they are not among a subject's strongest attitudes, as indexed by evaluation latency.

We had assumed that our preselected consistent attitude primes corresponded to the middle range of attitude assessment latencies for subjects in Experiment 1 on the basis of the response latencies of our normative subjects (see Table 2). To test this assumption directly, we computed for each Experiment 1 subject the mean latencies from the attitude assessment task for each of his or her idiosyncratic sets of fast-good, fast-bad, slow-good, and slow-bad attitude object primes, as well as for the four positively and four negatively evaluated attitude objects that served as that subject's consistent good and bad primes, respectively. These mean latencies were subjected to a 3 (prime type) × 2 (prime valence) ANOVA. Not surprisingly, there was a reliable main effect for prime type, $F(2, 44) = 109.98$, $p < .001$, $MS_e = 79{,}418.5$, which accounted for 65% of the total variance. Confirming the validity of our a priori set of consistent but

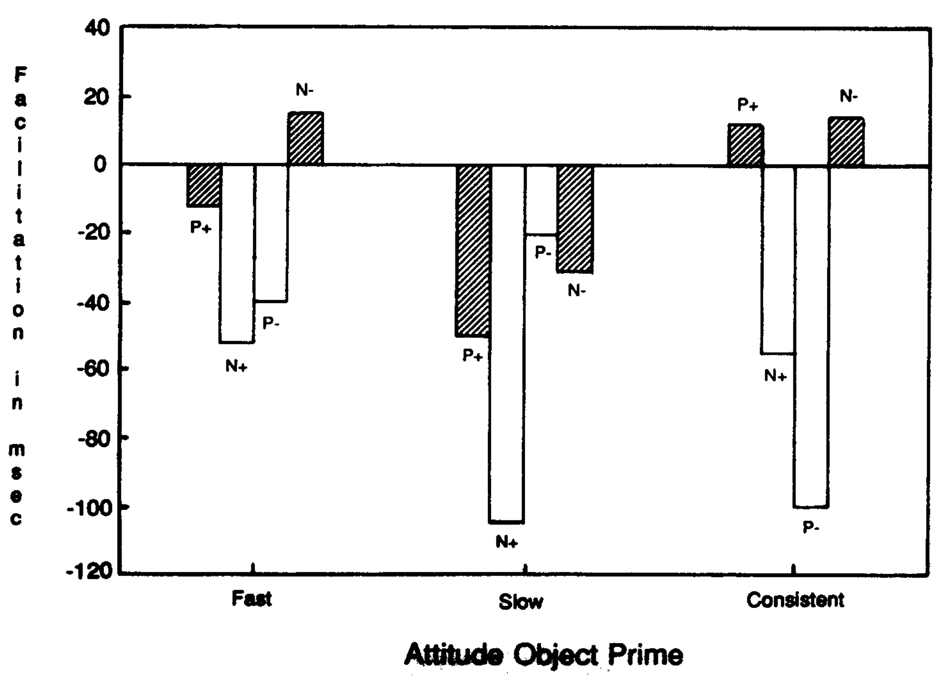

Figure 1. Mean target evaluation facilitation scores (in milliseconds) by prime type, prime valence (P = positive, N = negative), and target valence (+ = positive, − = negative), Experiment 1. (Mean baseline latencies: 755 ms for positive targets and 798 ms for negative targets.)

relatively slow primes, fast attitude object primes were evaluated in a mean time of 508 ms, slow primes in 1,354 ms, and consistent primes in 750 ms. Each of these means was reliably different from each other at $p < .001$. The mean evaluation time for the consistent primes was close to the grand mean of the entire distribution of all 2,116 evaluation latencies (23 subjects × 92 latencies each) of 769 ms (median = 693, $SD = 309$). The fast prime mean fell at the 10th percentile of the distribution (i.e., 10% of the evaluation latencies were faster), the consistent prime mean at the 59th percentile, and the slow prime mean at the 96th percentile.

Our results suggest that the automatic activation effect may be a fairly general phenomenon, as it holds for attitude objects distributed across most of the accessibility (evaluation latency) continuum (i.e., the consistent but slow attitude object primes; see Table 2). Before concluding that most attitudes are activated automatically in the natural social environment, however, it must first be determined that certain aspects of the automatic activation paradigm developed by Fazio et al. (1986) are not themselves responsible for producing the effect. Experiments 2 and 3 examined this issue. Experiment 2 was identical to Experiment 1 except that a 2-day delay between the assessment and priming tasks was introduced to remove any possible influence of recently thinking about one's attitude on the test of its automaticity. Experiment 3 also replicated Experiment 1's procedure, except that half of the subjects were not instructed to hold the prime (attitude object) word in memory while evaluating the target adjective.

Experiment 2

Category-priming studies have shown repeatedly that the processing of personality trait terms in one context causes the corresponding trait concepts to become more accessible and thus more likely to influence categorization and impressions in subsequent, ostensibly unrelated contexts (for reviews, see Higgins, 1989; Wyer & Srull, 1986). Other research has shown that people develop chronically accessible social constructs that are more accessible than other constructs even in the absence of any recent priming or preactivation; such constructs are automatically activated by the mere presence of relevant behavioral input (Bargh & Pratto, 1986; Bargh & Thein, 1985; Bargh & Tota, 1988; Higgins, King, & Mavin, 1982). Most important, the results of studies that have compared the effects of priming or temporary activation and the effects of automatic, chronically accessible constructs suggest that they are identical and interchangeable (Bargh et al., 1986; Bargh et al., 1988).

The implication of this research for the automatic activation of attitudes is that the automaticity apparently demonstrated in the Fazio et al. (1986) experiments and in our Experiment 1 may be conditional on subjects' having recently thought about their attitudes in the context of the initial attitude assessment task (see Bargh, 1989). Experiment 2 explored this possibility by having subjects complete the attitude assessment phase at a first experimental session and the priming task 2 days later at a second session. Thus, subjects performed the latter task, which assessed the automaticity of their attitudes, without having just been asked about their attitudes toward the objects.

Method

Subjects. Twenty-four students enrolled in the introductory psychology course at NYU participated in the experiment as partial fulfillment of a course requirement.

Materials, apparatus, and procedure. The stimulus materials, apparatus, and procedures were the same as in Experiment 1, except for the interpolation of a 2-day delay between the attitude assessment and priming tasks. After the subject had completed the attitude assessment task, the names of the attitude objects corresponding to his or her four fastest and four slowest good and bad responses were stored by the computer so that they could be used as priming stimuli in the second experimental session. The consistent priming stimuli were selected as in Experiment 1.

Results

As in Experiment 1, a latency in the priming task was excluded from the analyses if the corresponding evaluation of the adjective was incorrect (1.8% of the latencies were so discarded) or if it was less than 300 ms (1.1%),[5] and latencies over 2,500 ms (0.3% of latencies) were replaced by 2,500. The mean facilitation scores for each of the 14 experimental conditions were computed in the same manner as in Experiment 1, and the identical 3 (prime type) × 2 (prime valence) × 2 (target valence) ANOVA was computed. As in Experiment 1, the Prime Valence × Target Valence interaction was reliable, $F(1, 23) = 15.82$, $p < .001$, $MS_e = 7{,}824.5$, accounting for 3.6% of the variance. Simple main effects tests showed that evaluation latencies for positive targets were facilitated reliably more by positive than by negative primes, $F(1, 23) = 9.25$, $p < .01$, and those for negative targets were facilitated reliably more by negative than by positive primes, $F(1, 23) = 7.81$, $p < .02$.

However, in contrast with findings of Experiment 1, the Prime Type × Prime Valence × Target Valence interaction was not reliable, $F < 1$. As Figure 2 shows, the pattern of relative facilitation and inhibition that indicates the automatic attitude activation effect was present for all three prime types. The simple Prime Valence × Target Valence interaction was reliable in the fast, $F(1, 23) = 8.92$, and the consistent, $F(1, 23) = 6.33$ (both $ps < .01$), prime conditions and marginally reliable in the slow prime condition, $F(1, 23) = 2.71$, $p < .07$.

Discussion

Eliminating the immediately prior attitude assessment task from the paradigm did not eliminate the automatic activation effect for the fast and consistent attitude object primes. In fact, testing for automaticity of attitude activation without having the subject think about his or her attitude immediately beforehand resulted only in the reliable Prime Valence × Target Valence interaction that indicates the automaticity effect, with no reliable moderation of the effect depending on whether the prime was fast, consistent, or slow in evaluation speed.

The procedure of Experiment 2 was an improvement over that of earlier studies of the automatic attitude activation effect.

[5] We screened latencies for anticipations (i.e., less than 300 ms) in Experiment 1 as well, but there were none.

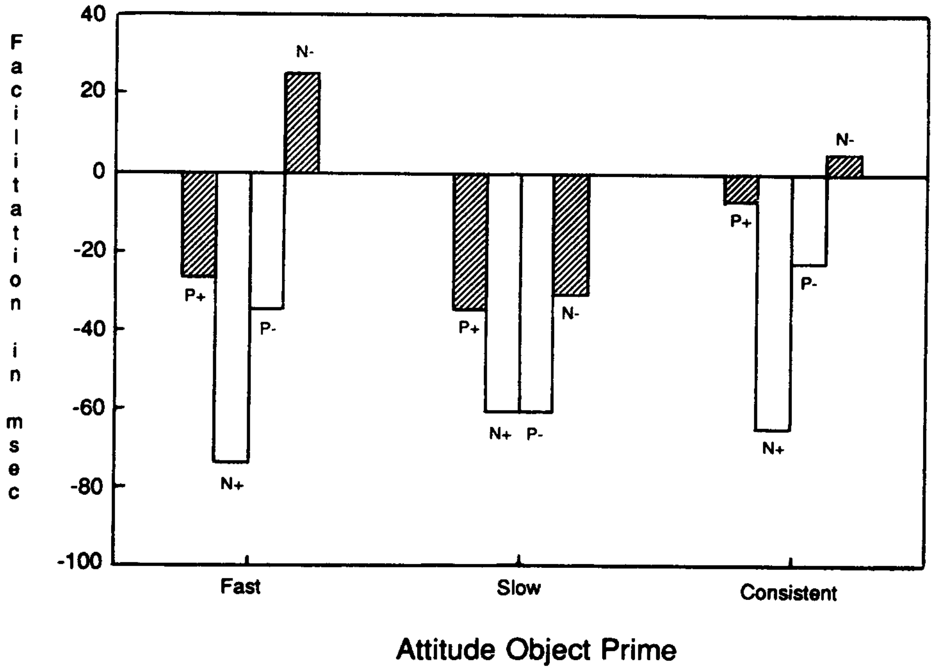

Figure 2. Mean target evaluation facilitation scores (in milliseconds) by prime type, prime valence (P = positive, N = negative), and target valence (+ = positive, − = negative), Experiment 2. (Mean baseline latencies: 669 ms for positive targets and 705 ms for negative targets.)

Removing the condition of having recently thought about one's attitude allows less ambiguous conclusions to be drawn about the effect of the mere presence of the attitude object. Therefore, in the more general case in which the person has not just thought about his or her attitude toward a given object, the automatic attitude activation effect is still obtained for attitude objects ranging across the continuum of associative strength.

Experiment 3

The procedure for the priming task in the Fazio et al. (1986) studies, and in our Experiments 1 and 2, calls for subjects to hold the attitude object prime in mind as a memory word while making their evaluative response to the target adjectives. This feature of the task may also be a contributor to the effect obtained. It may be the active, conscious rehearsal of the attitude object primes while the adjectives they were paired with were being evaluated that resulted in activation spreading from the attitude object to its associated evaluation. Consistent with this possibility, Warren (1977) found that facilitative priming effects across associative pathways increased in size as a function of how long the prime was presented (from 75 to 150 ms). Continuous conscious activation of the prime word representation throughout the trial therefore might have been more than sufficient to activate the associated evaluation as well. Without this intentional, conscious processing of the attitude object prime, the automatic attitude activation effect might not occur. At the least, this memory word aspect of the test of attitude automaticity must be examined before drawing the more general conclusion that the mere observation of an object can automatically elicit an evaluative response.

Method

Subjects. Fifty-nine NYU introductory psychology students participated in Experiment 3 as partial fulfillment of a course requirement.

Materials and apparatus. To provide an exact replication of the Fazio et al. (1986) Experiment 2 (300-ms SOA condition), we dropped the consistent prime stimuli used in our own Experiments 1 and 2. Thus, only fast and slow (good and bad) attitude objects served as primes. In all other respects, the materials and apparatus were the same as in those experiments.

Procedure. The procedure followed that used in Experiment 1, except that only half the subjects were instructed to hold the attitude object prime word in memory during each trial of the priming task. Subjects in the "no memory word" condition (n = 29) were told only that in the second task, there would be two words presented one after the other on the screen and that they were to press either the *good* or the *bad* button to indicate, as quickly as possible, whether the second word was positive or negative in meaning. No explanation or instructions were given to these subjects regarding the prime words. As in Experiment 1, subjects completed the priming task immediately after completing the attitude assessment task.

The rationale for the memory word instructions given by Fazio et al. (1986) was that otherwise subjects could adopt a strategy of ignoring the prime word on each trial because it was irrelevant to their task of evaluating the target word. To assess this possibility, we gave subjects 3 min immediately after the priming task to write down as many of the "first words" presented (i.e., primes) in the task as they could. By including this surprise memory test we could assess (albeit imperfectly;

see Bargh, 1984) whether subjects in the no memory word condition had attended to the priming stimuli.

Results

Prime recall. Subjects in both the memory word and no memory word conditions were able to recall the majority of priming stimuli presented to them. Of the 16 primes to which subjects were exposed, they recalled 9.7 (61%) of them on average. Not surprisingly, subjects instructed to remember the prime word on each trial recalled reliably more primes (M = 10.4, 65%) than did subjects not given memory instructions (M = 9.1, 57%); $t(57) = 2.09$, $p = .04$. Nonetheless, the high level of prime recall in the no memory instructions condition indicates that those subjects were attending to the prime presentation; the difference in recall is most likely due to extra processing (rehearsal) of the primes in the memory condition.

Automaticity test. Following the procedure of Experiments 1 and 2, priming task latencies were again excluded from analyses if the corresponding evaluation of the adjective was incorrect (1.6% of latencies) or if the latency was less than 300 ms (1.2% of latencies), and latencies over 2,500 ms were set to 2,500 (0.5% of latencies). As in the previous experiments, the mean facilitation scores for each of the 14 experimental conditions were computed and entered into a 2 (memory word instructions: yes vs. no) × 2 (prime type) × 2 (prime valence) × 2 (target valence) ANOVA. The Prime Valence × Target Valence interaction was again highly reliable, $F(1, 57) = 33.72$, $p < .001$, MS_e = 5,964.7, accounting for 4.4% of the total variance. As Figure 3 shows, the critical pattern of relative facilitation when prime and target valence match, and relative inhibition when the prime and target valence mismatch, obtained for both fast and slow prime types.

This two-way interaction was qualified, however, by a reliable three-way interaction involving prime type, $F(1, 57)$ = 12.37, $p = .001$, MS_e = 4,001.0 (1.1% of total variance). Within this three-way interaction, simple effects tests revealed that the simple Prime Valence × Target Valence interaction was reliable for fast primes, $F(1, 57) = 56.98$, $p < .001$, and also for slow primes, $F(1, 57) = 6.32$, $p = .02$. Thus, the three-way interaction can be attributed to the stronger effect obtained for the fast than for the slow prime condition, although the effect per se held true regardless of prime type. None of the main effects or interactions involving the between-subjects factor of memory word instructions proved reliable.[6]

Discussion

The results of Experiment 3 indicate that the automatic attitude activation effect is not dependent on the memory word instruction feature of the original paradigm. Thus, when the requirement for subjects to hold the attitude object prime in working memory during each trial of the automaticity test is removed, providing a better test of the "mere presence" hypothesis, the automatic activation effect is still obtained. Moreover, the effect is obtained for the subject's slowest (weakest) as well as his or her fastest (strongest) attitudes. Thus, Experiment 3 is another demonstration that under conditions more closely approximating the mere presence of the object in the environment, the effect occurs for a majority of objects for which one has a stored evaluation. Finally, the reliable three-way interaction, in which the effect was found to be statistically stronger for fast than for slow attitude object primes, replicated the findings of Experiment 1, in which the attitude assessment phase also came immediately before the test of automaticity.

Overall Regression Analysis

We began this line of research by collecting normative data on characteristics of attitude object stimuli that might covary with speed of object evaluation or associative strength. As shown in Table 1, all of these factors correlated reliably with mean evaluation latency. However, a correlation with mean attitude object evaluation latency is not the same as a correlation with the automatic activation effect itself. Therefore, using the data from Experiments 1–3, we examined the extent to which these factors moderated the automaticity effect; that is, the signature Prime Valence × Target Valence interaction.

To do this we conducted a hierarchical multiple regression analysis, in which we attempted to predict each response latency in the automaticity task for each subject in each experiment. This entailed creating a data file that included all relevant information about the attitude object prime and the target adjective for that trial, both in normative terms (across all subjects for which we had that information) as well as idiosyncratically for the given subject. These data included the following: normative prime valence (dichotomous good vs. bad classification based on the normative data from our mass testing session; see Appendix), normative prime evaluation latency (across all subjects in the attitude assessment tasks of Experiments 1–3 as well as of the normative consistency and ambivalence experiments), idiosyncratic prime evaluation latency (the subject's own evaluation latency from the attitude assessment task), normative extremity (absolute value of mean evaluation from mass testing session), normative ambivalence and normative polarization (from normative ambivalence study; see Appendix), idiosyncratic target evaluation (positive or negative evaluation of target on that trial), idiosyncratic mean response latency (the subject's mean response latency over all trials in the automaticity task), and normative prime evaluation consensus index.

This last predictor was based on the attitude object evaluations of all subjects in Experiments 1–3 and the normative studies (i.e., the mass testing session, the consistency experiment, and the ambivalence experiment)—a total sample size of 412, giving a more stable estimate of consistency than did the earlier estimate based on the 30 normative consistency experiment subjects. We created a consensus factor score comprised of two elements: the proportion of the 412 respondents who gave the

[6] There was also a main effect for target valence, $F(1, 57) = 24.74$, $p <$ .001 (9.6% of variance), which was qualified by a Target Valence × Prime Type interaction, $F(1, 57) = 11.57$, $p = .002$ (0.7% of variance). As can be seen in Figure 3, the general tendency for negative targets to have faster evaluation times than positive targets was stronger in the slow prime condition. All other main effects and interactions were nonsignificant at $p > .20$.

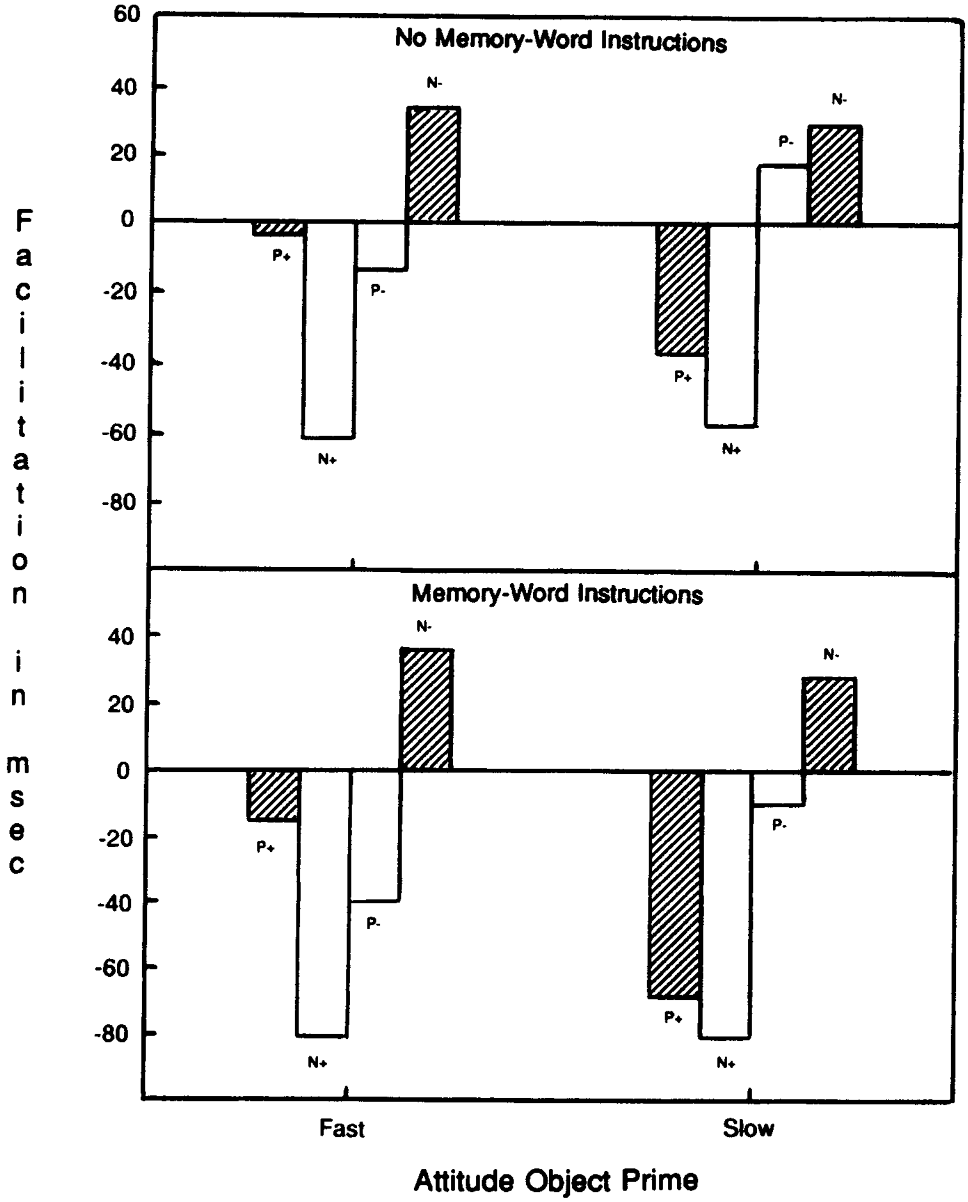

Figure 3. Mean target evaluation facilitation scores (in milliseconds) by prime type, prime valence (P = positive, N = negative), and target valence (+ = positive, − = negative), separately for the no memory word and the memory word instructions conditions, Experiment 3. (Mean baseline latencies: for no memory word condition, 529 ms for positive targets and 588 ms for negative targets; for memory word condition, 626 ms for positive targets and 709 ms for negative targets.)

modal evaluation (i.e., good vs. bad) and the inverse of the standard deviation of the evaluative ratings of the attitude objects from the mass testing session (n = 274). The former number reflects the amount of consensus as to the valence of the object; the latter number represents the amount of variability in the evaluation of the object on the questionnaire. As they were highly correlated across the 92 attitude objects, r = .76, the mean of the two scores served as our evaluative consensus index in the regression analysis (these scores are also presented in the Appendix). Although this consensus index can be interpreted as a proxy for consistency within the individual subject, it would nonetheless be desirable for future research to more directly assess intraindividual attitude consistency (as well as intraindividual ambivalence, polarization, and extremity) as a moderator of automatic attitude activation.[7]

[7] It is noteworthy that this composite consensus score itself correlated highly with the measure we used to select the consistent primes for Experiments 1 and 2 (r = .72).

The setwise hierarchical multiple regression (see Cohen & Cohen, 1975, Chapter 4) proceeded by regressing the subject's response latency onto these predictor variables for each trial.[8] Excluding the baseline prime trials, each subject in Experiments 1 and 2 contributed 168 response latencies to the data file, and each subject in Experiment 3 contributed 80.[9] There were 3,024 (18 × 168) trials from Experiment 1, 1,848 (11 × 168) from Experiment 2, and 4,720 (59 × 80) from Experiment 3 to be potentially included in the regression data file, a total of 9,592. After we eliminated those latencies below 300 ms and also incorrect responses on the adjective evaluation task (and, as before, setting all latencies greater than 2,500 equal to 2,500), a final set of 8,882 latencies remained.

Table 3 presents a summary of the regression analysis. The first step was to extract the variance in individual trial response latencies due to the subject's idiosyncratic speed of responding, and so the subject's mean response latency (MRT) on the adjective evaluation task was entered as a criterion-coded variable.[10] Next, the valence of the target adjective on that trial was entered, coded as either −1 (negative) or 1 (positive), followed by the normative prime valence in Step 3. In Step 4, the interactions between subject's MRT and target valence and between subject's MRT and prime valence were entered to extract the variance attributable to the subject being faster to respond to one prime or target valence than the other.

The fifth step, in which the interaction between normative prime valence and target valence was entered, is critical as it provides a general test of the existence of the automatic attitude activation effect across all three experiments. It was indeed a reliable predictor of response latencies. Prime–target valence matches (either −1 or 1 squared, so that the interaction variable equaled 1) were associated with faster response latencies than were mismatches (interaction variable equal to −1).

In Step 6, the variance due to the three-way interaction of this factor with subject's MRT was extracted. At this point in the analysis, all main effects and interactions other than the critical predictors of the automaticity effect had been entered, and their associated variance removed. Thus, in Step 7 we entered the normative consensus index, normative prime evaluation latency, idiosyncratic prime evaluation latency, normative attitude extremity, normative ambivalence, and normative polarization scores. This set the stage for the critical Step 8, in which the interactions of these variables with the automaticity effect (the Prime Valence × Target Valence interaction) were entered.

Because of the high intercorrelations among the predictor variables,[11] interpretation of the Step 8 results was potentially problematic. If all predictors are entered simultaneously, the results could be contaminated by multicollinearity (see Cohen & Cohen, 1975). If each of the predictors is entered separately and singly, however, a significant result could be due to that variable's relation to another predictor. Consequently, we first entered the Step 8 predictor variables individually to rule out potential multicollinearity distortions. Then we checked those results against a simultaneous-entry Step 8 to discover the independent effects of each variable when all others are statistically controlled.

The results of the individual-entry Step 8 are presented in Table 3. This step of the analysis revealed several attitude object characteristics that moderated the automatic activation effect. The greater the normative ambivalence for an attitude object prime, the weaker the automaticity effect.[12] Moreover, normative evaluation latency also reliably covaried with the automaticity effect; the smaller the latency, the greater the effect. However, the subject's idiosyncratic prime evaluation latency—the attitude strength variable that was the focus of the Fazio et al. (1986) research—did not predict the automaticity effect. Normative polarization was marginally related to the effect (p = .08). Finally, the follow-up regression analysis with simultaneous entry of the Step 8 predictors confirmed that ambivalence and normative evaluation latency continued to be reliable moderators of the automaticity effect when the effects of the other predictors were partialed out.

One aspect of the overall regression analysis deserves comment. We used normative prime valence as our index of whether the prime on each trial was positive or negative, instead of the subject's idiosyncratic evaluation of the prime in the attitude assessment phase of the experiment. Subjects' own prime evaluations differed from the normative evaluation on 1,321 out of the 8,882 trials, or 15% of the time. We chose the normative index because of its greater stability and because many of the cases on which the two measures differed appeared to be errors by the subjects (e.g., *recession* was frequently evaluated as good).

Nonetheless, we recomputed the basic regression analysis, using only those cases in which the normative and idiosyncratic prime valence indexes were the same. With individual entry of the Step 8 predictors, only normative prime evaluation latency covaried with the automaticity effect, $t = 1.74$, $p = .08$; for all other predictors, $ts < 1.10$. In the simultaneous analysis, normative evaluation latency was reliable ($t = 2.20$, $p < .03$), and ambivalence proved marginally reliable ($t = 1.65$, $p < .10$).

Next, we reran the regression analysis on all cases in the

[8] Tony Greenwald and an anonymous reviewer suggested this analysis to us.

[9] It was not possible to recover the information as to which attitude objects served as primes for 5 subjects in Experiment 1 and for 13 subjects in Experiment 2; these words were never part of the computer-written data file.

[10] This predictor accounted for nearly all of the explained variance in the regression, $t = 71.71$, $p < .0001$, $R^2 = .36$. The R^2 for the complete model was only .37. Although the remaining predictors account for little of the variability in response latencies, it should be remembered that the focus of the research is not to predict latencies, but to determine the conditions under which the attitude automaticity effect occurs. In other words, the size of the effect on latencies is not important; what is critical is the sheer existence of the effect and its generality.

[11] The correlations among the normative predictors can be found in Table 1. Idiosyncratic prime evaluation latency also correlated moderately but reliably (all $ps < .001$) with the normative predictors across all trials in the regression data set (rs = −.36 with consensus, .47 with normative latency, −.34 with extremity, .12 with ambivalence, and −.23 with polarization).

[12] Note that, as the automaticity effect (Prime Valence × Target Valence interaction) is itself negative, three-way interactions involving it that have an associated positive t value indicate an attenuation of the effect with increasing scores on that variable, whereas negative t values indicate that increasing scores on the interacting variable are related to an enhancement of the effect.

Table 3
Summary of the Overall Regression Analysis

Sequence	Regression weight	Standard error	t value
Step 1			
Subject's mean response latency (MRT)	0.991	0.014	71.69***
Step 2			
Target adjective valence (TAV)	−3.510	2.771	−1.27
Step 3			
Normative prime valence (NPV)	−0.861	2.790	<1
Step 4			
MRT × TAV	−0.030	0.014	−2.20*
MRT × NPV	0.029	0.014	2.06*
Step 5 (Automaticity effect)			
TAV × NPV (AUTO)	−14.900	2.790	−5.35***
Step 6			
MRT × AUTO	−0.036	0.014	−2.58**
Step 7			
Normative consensus (NC)	−5.505	67.485	<1
Normative prime evaluation latency (NPRT)	0.115	0.044	2.59**
Idiosyncratic prime evaluation latency (IPRT)	0.009	0.006	1.50
Normative extremity (NEX)	6.108	9.290	<1
Normative ambivalence (NA)	−3.524	13.543	<1
Normative polarization (NP)	−4.630	20.264	<1
Step 8 (entered individually)			
NC × AUTO	−1.103	21.847	<1
NPRT × AUTO	0.061	0.028	2.15*
IPRT × AUTO	0.004	0.006	<1
NEX × AUTO	−2.510	2.389	−1.05
NA × AUTO	11.554	4.930	2.34*
NP × AUTO	−8.934	5.100	−1.75

* $p < .05$. ** $p < .01$. *** $p < .001$.

original analysis using idiosyncratic prime evaluation in place of normative evaluation. With individual entry at Step 8, both normative prime evaluation latency (t = 2.24, p < .025) and idiosyncratic prime evaluation latency (t = 2.13, p < .04) were reliable (all other ps > .25). As noted earlier, however, significance of a predictor when entered by itself might be caused by its relation to another predictor. When the independent influences of normative and idiosyncratic evaluation latencies were assessed in the simultaneous-entry version of the analysis, normative evaluation latency remained (marginally) reliable (t = 1.83, p < .07), whereas idiosyncratic evaluation latency did not (t = 1.30, p = .20), confirming the outcome of the original regression analysis.

As indicated by the regression results, the automaticity effect varied in strength primarily as a function of normative—but not idiosyncratic—attitude object evaluation latency and to a certain extent as a function of normative ambivalence. Tentatively, these findings suggest that the strength of the automatic activation effect is a function not of variations in the accessibility of the individual subjects' attitudes toward the object but of features of the object representation or its evaluation that are constant across individuals. We return to this point later.

General Discussion

Fazio et al. (1986) demonstrated that attitudes corresponding to a subject's 8 most quickly evaluated attitude objects among a set of 92 were capable of becoming automatically active, whereas attitudes toward the 8 most slowly evaluated attitude objects were not. On the basis of their assumption that evaluation speed was an index of accessibility and its major determinant, associative strength, they concluded that one's strongest, most accessible attitudes became active automatically on the mere presence of the attitude object.

Our concern with the generality of this important finding sprang both from a lack of information about the ability of the majority of attitude objects in that study to show the effect—namely, the middle range of attitude objects with evaluation latencies falling between the eight fastest and eight slowest—and from the results of Fazio et al.'s (1986) Experiment 3, in which the automatic attitude evaluation effect was obtained for relatively slowly but consensually evaluated attitude objects. On the basis of our normative data collection, we selected attitude objects that were consistently but relatively slowly evaluated across subjects for inclusion as consistent primes in Experiment 1. That study, which otherwise followed the Fazio et al. (1986) Experiment 2 procedure exactly, replicated the previous results by obtaining the automatic activation effect for the subject's idiosyncratically selected fast but not slow primes. However, the effect was also obtained for the preselected consistent but relatively slowly evaluated attitude object primes. In other words, the automatic attitude activation effect appeared to occur for the majority of the attitude object stimuli in the Fazio et al. (1986) paradigm; the exceptional cases were the subject's slowest evaluated attitude objects that did not reliably show the effect.

Whereas Experiment 1 was focused on the internal generalizability of the automatic activation effect (i.e., across the range of attitude objects within the original paradigm), Experiments 2 and 3 were concerned with whether the effect would occur when the experimental situation was modified to more closely approximate conditions of the mere presence of the attitude object in the natural environment. Experiment 2 separated in time the attitude assessment phase from the priming task that assessed attitude automaticity, so that the immediately prior attitude assessment task could not itself make the attitude object evaluations temporarily accessible, thus simulating automaticity for some time thereafter (see Bargh, 1989, 1992; Spielman & Bargh, 1991). Somewhat surprisingly, the effect of the interpolated delay was not to eliminate the effect for the fast and consistent primes; instead, the signature Prime Valence × Target Valence interaction was itself reliable, and its size did not fluctuate reliably as prime type varied from fast to consistent to slow.

Experiment 3 was also a replication of the basic paradigm, this time with the alteration of removing the memory word instructions for half of the subjects. By having subjects hold the attitude object prime in memory until they had evaluated the adjective target on each trial in the priming task, the original paradigm did not permit unambiguous conclusions about whether the automatic activation effect would occur without such deliberate, intensive conscious thought about the attitude object. Of course, if the effect required only the mere presence of the attitude object to occur, such memory word instructions should not be necessary to produce the effect in the laboratory. The results of Experiment 3 showed that indeed they were not necessary; the automatic activation effect held for both the subject's fastest and slowest evaluated attitude objects.

We conducted further statistical tests of whether the automaticity effect (i.e., the simple Prime Valence × Target Valence interaction) held reliably across the three experiments for the subjects' slow and fast attitude object primes as well as for the consistent but slow primes used in Experiments 1 and 2.[13] To do so we used meta-analytic procedures for combining results across studies and for assessing differences in the size of the effect across studies (Rosenthal, 1978; Rosenthal & Rubin, 1979). Not surprisingly, given its reliability in each of the three studies individually, the automaticity effect for the fast attitude objects was reliable overall, with weighted average $Z = 4.42$, $p < .001$, average effect size = .43; moreover, the effect sizes did not differ across the three experiments, $\chi^2(2, N = 106) = 0.29$, $p = .86$. Consistent attitude objects presented to subjects in Experiments 1 and 2 (see Table 2) also showed a reliable automaticity effect, $Z = 4.11$, $p < .001$, average effect size = .60, which did not vary across the two studies, $\chi^2(1, N = 47) = .06$, $p = .81$. Of greatest importance, however, the same pattern of results was obtained in the case of the slow attitude object primes. The automaticity effect was reliable, $Z = 2.66$, $p < .01$, effect size = .26, and it did not vary reliably across the three experiments, $\chi^2(2, N = 106) = 0.52$, $p = .77$. Finally, comparisons between prime types showed that the automaticity effect was reliably greater for fast than slow primes ($p < .004$) and also reliably greater for consistent than slow primes ($p < .03$). However, the automaticity effect proved nonreliably smaller for fast than consistent primes ($Z < 1$), even though the latter primes were associated with reliably slower response times than the fast primes.

Taken together, the results of the three experiments suggest the automatic attitude activation effect is quite general, holding across most if not all of the range of 92 attitude objects that served as stimuli. These attitude objects varied widely as to their extremity, ambivalence, and polarization of attitude, their consistency of evaluation across subjects (i.e., consensus), and their mean evaluation latencies (see Appendix). Moreover, removing features of the paradigm that potentially could have contributed to the automaticity effect did not change its strength across the three experiments, demonstrating its relatively unconditional nature (see Bargh, 1989).

Implications for Understanding the Concept of Attitude Strength

Our meta-analysis of the automaticity effect across the three experiments showed it to be reliable for attitudes at the extreme slow tail of the idiosyncratic evaluation latency distribution, as well as for attitudes at the extreme fast tail. If, as postulated by the attitude accessibility model, idiosyncratic evaluation latency for a given attitude object is an index of the accessibility or associative strength of the corresponding attitude, then our results indicate either that (a) attitude accessibility or associative strength is irrelevant to the occurrence of automatic attitude activation or (b) the threshold on the associative strength continuum for the occurrence of the effect is quite low. The present results stand in contrast with theoretical statements of the attitude accessibility model, which imply that relatively high associative strength is a prerequisite for automatic activation (see the introduction).

Although we have focused exclusively on the automatic attitude activation effect, our results also have implications for other research guided by the attitude accessibility model. As noted at the outset of this article, associative strength is the central construct in the attitude accessibility model's description of the process by which attitudes guide behavior (Fazio, 1986, 1989, 1990). Thus, associative strength is viewed as the major determinant of an attitude's chronic accessibility, and it is the chronic accessibility of an attitude that "determines the power and functionality of an attitude" (Fazio, 1989, p. 154). According to the model, highly accessible (vs. less accessible) attitudes are more likely to guide the processing of attitude-relevant information and, hence, more likely to guide behavior. Thus, stronger selective perception effects as well as stronger attitude–behavior correlations are predicted when attitude accessibility (i.e., associative strength) is higher (vs. lower).

[13] It is important to keep in mind that whereas we refer here to fast and slow sets of attitude objects—because they were selected as primes on the basis of subjects' idiosyncratic evaluation latencies—these sets also differ in terms of normative evaluation latency, ambivalence, and other attitude object characteristics (see Table 1). Thus, any differences in the size of the automaticity effect as a function of prime type cannot be unequivocally attributed to differences in individual attitude accessibility.

Numerous studies by Fazio and his colleagues have supported these predictions of the model (see Fazio, 1989, 1990 for reviews). All of this research, however, has relied on evaluation latency to assess individual differences in associative strength or accessibility or to assess whether variables presumed to affect associative strength (e.g., direct experience with attitude objects; repeatedly expressing an attitude) in fact did so. For example, Fazio and Williams (1986) concluded that associative strength moderated the predictability of behavior from attitudes on the basis of data showing higher attitude–behavior correlations among subjects who responded more quickly (vs. less quickly) to an attitude query. However, our finding of reliable correlations between evaluation latency and a number of attitudinal qualities (e.g., ambivalence, consistency, and extremity; see Table 1) and a reliable moderation of the automaticity effect by normative (but not idiosyncratic) evaluation latency in the regression analyses suggest that such conclusions may be premature.

Specifically, our findings call into question the assumption that the variety of variables that researchers have identified as indicators of attitude strength are reducible to a single construct that is best conceptualized in terms of the associative connection between the attitude object and a single evaluation. Ambivalence did seem to moderate the automaticity effect but is not well represented in terms of this single evaluation model (i.e., either a good or a bad evaluation), because ambivalence refers to the relative strengths of both good and bad feelings toward the object (Garner, 1962; Kaplan, 1972). Of course, given the high intercorrelations shown in Table 1 among the various predictors of attitude object evaluation latency, any conclusions about the central importance of any one predictor is somewhat speculative. Further experimentation on the role of attitude ambivalence in moderating the automatic activation effect is plainly called for. The important point for the present discussion is that attitude features such as ambivalence, polarization, extremity, and normative evaluation latency are predictive of evaluation latency and may moderate the automaticity effect, yet these additional indicators are not necessarily reducible to the same construct that presumably underlies accessibility—the strength of the association in memory between the object representation and its evaluation.

It is intriguing that normative attitude object evaluation latency reliably covaried with the automaticity effect. This finding suggests that there is some characteristic of the attitude object itself (or at least of our shared cultural experience of it; see following discussion) independent of the accessibility of the individual subject's attitude toward it that increases the strength of the automatic activation effect.

Just what that something is remains unclear.[14] In our normative study we found that evaluation extremity, consistency, ambivalence, and polarization all were reliably correlated with normative evaluation latency (see Table 1). Yet these factors were included as separate predictors in the regression analysis, and normative evaluation latency continued to be an independent significant predictor of the automaticity effect. Because normative evaluation latency seemed to be the primary moderator of automaticity, perhaps individual variations in attitude strength matter less to the occurrence of the effect than do commonalities in how people react to and feel about the attitude object. These may be culturally transmitted (e.g., anti-Semitism or pro-individualism) or attributable more directly to qualities of the attitude object (e.g., the beauty of a rose or the obnoxiousness of car alarms). Whatever the cause or source, the implication of our findings is that, at least within a given culture, the same attitudes will be more strongly automatically activated for most people, and the same ones will be more weakly automatically activated. Another way of saying this is that the variability in the automatic activation effect may be among attitude objects and not among people—that the source of the effect lies in the environmental stimuli and not the individual perceiver (e.g., McArthur & Baron, 1983).

Is (Nearly) Everything Preconsciously Evaluated?

Beyond its consequences for existing models of the attitude–behavior relation, the finding that most attitude object stimuli are capable of automatically activating their associated evaluations in memory has more general ramifications for the nature and extent of preconscious information processing, for the affect–cognition interface, and for the way in which evaluatively toned information is organized in memory. If the majority of environmental stimuli for which one has a stored evaluation activate that evaluation automatically on their mere presence, what are the consequences of the activated evaluation for judgment and behavior? Is this general automatic evaluation effect a precursor to emotional or mood states (cf. Lazarus, 1991)?

Such questions must await the outcome of further, rigorous investigation of just how general the automatic evaluation effect is. A thorough assessment of the generality of the automatic attitude activation effect must go beyond the particular attitude objects and the particular paradigm that we and Fazio et al. (1986) have used to test for the presence of the effect. With regard to the generality of this particular set of 92 attitude objects, however, the wide variety of social and nonsocial objects and concepts represented—activities, states of being, foods, and famous people—and the wide range of evaluative reactions that subjects have to them lead us to conclude that the automatic activation effect does generalize well across varieties of evaluative stimuli. Fazio et al. (1986) had selected the set of 92 attitude objects to be widely representative and concluded that their findings were "of relevance to any broadly defined 'object' towards which an individual possesses some affective linkage" (p. 236). A cursory examination of the 92 attitude object stimuli in the present Appendix shows them to be rather innocuous and mundane for the most part; not the names of objects toward which fiery passions would be ignited (e.g., abortion, death penalty, gun control, and school prayer). If the effect occurs for the majority of these objects, it is highly probable that it occurs for more affectively charged attitudes as well.

[14] It is not the frequency or length of the attitude object prime word. We reran the basic regression analysis shown in Table 3, including the main effects of these factors at Step 7 and their interactions with the automaticity effect at Step 8. Neither did any factor interact with the effect ($ts < 1$) nor did their inclusion in the analysis alter the reliability or sign of any other factor.

References

Ajzen, I., & Fishbein, M. (1980). *Understanding attitudes and predicting social behavior.* Englewood Cliffs, NJ: Prentice-Hall.

Anderson, N. H. (1968). Likableness ratings of 555 personality trait words. *Journal of Personality and Social Psychology, 9,* 272–279.

Bargh, J. A. (1984). Automatic and conscious processing of social information. In R. S. Wyer, Jr., & T. K. Srull (Eds.), *Handbook of social cognition* (Vol. 3, pp. 1–43). Hillsdale, NJ: Erlbaum.

Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influence in social perception and cognition. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 3–51). New York: Guilford Press.

Bargh, J. A. (1992). The ecology of automaticity: Towards specifying the conditions necessary to produce automatic processing effects. *American Journal of Psychology, 105,* 181–199.

Bargh, J. A., Bond, R. N., Lombardi, W. J., & Tota, M. E. (1986). The additive nature of chronic and temporary sources of construct accessibility. *Journal of Personality and Social Psychology, 50,* 869–878.

Bargh, J. A., Lombardi, W. J., & Higgins, E. T. (1988). Automaticity of chronically accessible constructs in Person × Situation effects on person perception: It's just a matter of time. *Journal of Personality and Social Psychology, 55,* 599–605.

Bargh, J. A., & Pratto, F. (1986). Individual construct accessibility and perceptual selection. *Journal of Experimental Social Psychology, 22,* 293–311.

Bargh, J. A., & Thein, R. D. (1985). Individual construct accessibility, person memory, and the recall-judgment link: The case of information overload. *Journal of Personality and Social Psychology, 49,* 1129–1146.

Bargh, J. A., & Tota, M. E. (1988). Context-dependent automatic processing in depression: Accessibility of negative constructs with regard to self but not others. *Journal of Personality and Social Psychology, 54,* 925–939.

Cohen, J., & Cohen, P. (1975). *Applied multiple regression/correlation analysis for the behavioral sciences.* Hillsdale, NJ: Erlbaum.

Converse, P. E. (1970). Attitudes and non-attitudes: Continuation of a dialogue. In D. E. Apter (Ed.), *Ideology and discontent* (pp. 206–261). New York: Free Press.

Doll, J., & Ajzen, I. (1992). *Accessibility and stability of predictors in the theory of planned behavior.* Unpublished manuscript, University of Massachusetts, Amherst.

Downing, J. W., Judd, C. M., & Brauer, M. (in press). Effects of repeated expression on attitude extremity. *Journal of Personality and Social Psychology.*

Eagly, A., & Chaiken, S. (in press). *The psychology of attitudes.* San Diego, CA: Harcourt Brace & Jovanovich.

Fazio, R. H. (1986). How do attitudes guide behavior? In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition* (Vol. 1, pp. 204–243). New York: Guilford Press.

Fazio, R. H. (1989). On the power and functionality of attitudes: The role of attitude accessibility. In A. R. Pratkanis, S. J. Breckler, & A. G. Greenwald (Eds.), *Attitude structure and function* (pp. 153–179). Hillsdale, NJ: Erlbaum.

Fazio, R. H. (1990). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. In M. P. Zanna (Ed.), *Advances in experimental psychology* (Vol. 23, pp. 75–109). San Diego, CA: Academic Press.

Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology, 50,* 229–238.

Fazio, R. H., & Williams, C. J. (1986). Attitude accessibility as a moderator of the attitude–perception and attitude–behavior relations: An investigation of the 1984 presidential election. *Journal of Personality and Social Psychology, 51,* 505–514.

Fisk, A. D., & Schneider, W. (1984). Consistent attending versus consistent responding in visual search: Task versus component consistency in automatic processing development. *Bulletin of the Psychonomic Society, 22,* 330–332.

Francis, W. N., & Kucera, H. (1982). *Frequency analysis of English usage: Lexicon and grammar.* Boston, MA: Houghton Mifflin.

Garner, W. R. (1962). *Uncertainty and structure as psychological concepts.* New York: Wiley.

Higgins, E. T. (1989). Knowledge accessibility and activation: Subjectivity and suffering from unconscious sources. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 75–123). New York: Guilford Press.

Higgins, E. T., King, G. A., & Mavin, G. H. (1982). Individual construct accessibility and subjective impressions and recall. *Journal of Personality and Social Psychology, 43,* 35–47.

Judd, C. M., & Kulik, J. A. (1980). Schematic effects of social attitudes on information processing and recall. *Journal of Personality and Social Psychology, 46,* 1193–1207.

Kaplan, K. J. (1972). On the ambivalence–indifference problem in attitude theory and measurement: A suggested modification of the semantic differential technique. *Psychological Bulletin, 77,* 361–372.

Katz, I. (1981). *Stigma: A social psychological analysis.* Hillsdale, NJ: Erlbaum.

Lazarus, R. S. (1991). *Emotion and adaptation.* New York: Oxford University Press.

Logan, G. D. (1980). Attention and automaticity in Stroop and priming tasks: Theory and data. *Cognitive Psychology, 12,* 523–553.

Logan, G. D. (1989). Automaticity and cognitive control. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 52–74). New York: Guilford Press.

McArthur, L. Z., & Baron, R. M. (1983). Toward an ecological theory of social perception. *Psychological Review, 90,* 215–238.

McCloskey, M. E., & Glucksberg, S. (1978). Natural categories: Well defined or fuzzy sets? *Memory & Cognition, 6,* 462–472.

Milner, P. M. (1957). The cell assembly: Mark II. *Psychological Review, 64,* 242–252.

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.

Posner, M. I. (1978). *Chronometric explorations of mind.* Hillsdale, NJ: Erlbaum.

Posner, M. I., & Snyder, C. R. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola symposium* (pp. 55–85). Hillsdale, NJ: Erlbaum.

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

Rosenthal, R. (1978). Combining results of independent studies. *Psychological Bulletin, 86,* 1165–1168.

Rosenthal, R., & Rubin, D. B. (1979). Comparing significance levels of independent studies. *Psychological Bulletin, 85,* 185–193.

Sanbonmatsu, D. M., Osborne, R. E., & Fazio, R. H. (1986, May). *The measurement of automatic attitude activation.* Paper presented at the annual meeting of the Midwestern Psychological Association, Chicago.

Schneider, W., & Fisk, A. D. (1982). Degree of consistent training: Improvements in search performance and automatic process development. *Perception & Psychophysics, 31,* 160–168.

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

Shallice, T. (1972). Dual functions of consciousness. *Psychological Review, 79,* 383–393.

Spielman, L. A., & Bargh, J. A. (1991). *The influence of immediately prior administration of depression inventories on experimental tests of depressive cognition: The "BDI-Priming Effect."* Unpublished manuscript, New York University.

Warren, R. E. (1977). Time and the spread of activation in memory. *Journal of Experimental Psychology: Human Learning and Memory, 3,* 458–466.

Whaley, C. P. (1978). Word-nonword classification time. *Journal of Verbal Learning and Verbal Behavior, 17,* 143–154.

Wyer, R. S., Jr., & Gordon, S. E. (1984). The cognitive representation of social information. In R. S. Wyer, Jr., & T. K. Srull (Eds.), *Handbook of social cognition* (Vol. 2, pp. 73–150). Hillsdale, NJ: Erlbaum.

Wyer, R. S., Jr., & Srull, T. K. (1986). Human cognition in its social context. *Psychological Review, 93,* 322–359.

Appendix

Characteristics of the 92 Attitude Objects

Attitude object	Frequency[a]	Length[b]	Evaluation[c]	Consistency[d]	Consensus[e]	Ambivalence[f]	Polarization[g]	Latency[h]
Alcoholism	1	10	−4.1	100	.81	0.4	2.7	806
Anchovies	1	9	−1.8	80	.50	0.7	1.7	1,010
Aquarium	1	8	1.9	90	.66	0.9	1.7	813
Baby	57	4	3.5	97	.72	1.5	1.7	675
Basketball	9	10	1.2	87	.53	1.3	1.2	807
Beer	36	4	−0.5	87	.40	1.4	1.4	822
Birthday	18	8	3.0	100	.69	1.3	1.7	646
Bombs	68	5	−4.0	93	.76	0.5	2.6	704
Butterfly	3	9	2.5	97	.70	0.6	2.0	709
Cake	16	4	2.4	100	.71	1.8	1.4	681
Cancer	24	6	−4.4	97	.80	0.0	3.0	674
Cavities	13	8	−3.2	93	.71	0.1	2.7	827
Chocolate	9	9	2.8	97	.67	2.1	1.3	701
Circus	7	6	1.9	97	.66	1.1	1.7	656
Clothes	89	7	3.2	97	.76	1.2	1.8	766
Clown	5	5	1.7	90	.63	0.9	1.5	736
Cockroach	2	9	−3.8	97	.79	0.2	2.6	737
Coffee	78	6	0.3	87	.44	1.6	1.3	818
Crime	49	5	−3.9	93	.80	0.1	2.8	709
Dancing	13	7	3.1	100	.70	0.3	2.3	661
Death	284	5	−3.1	97	.61	0.9	2.1	724
Dentist	16	7	−1.1	77	.46	2.3	1.0	892
Disco	1	5	0.6	87	.45	1.8	1.1	864
Disease	72	7	−3.9	100	.80	0.3	2.8	708
Divorce	23	7	−2.6	97	.64	1.3	1.8	735
Dormitory	6	9	1.1	73	.52	2.1	0.8	1,014
Eagle	12	5	1.9	87	.57	0.8	1.5	754
Exams	38	5	−1.2	77	.51	2.1	1.3	896
Flowers	56	7	3.5	100	.77	0.6	2.2	671
Food	198	4	3.0	93	.74	1.8	1.8	660
Football	37	8	0.9	90	.49	1.6	1.0	757
Fraternity	7	10	0.0	83	.44	2.0	0.9	1,027
Friday	64	6	3.8	83	.77	0.8	2.3	795
Friend	294	6	4.4	100	.89	1.0	2.2	669
Funeral	31	7	−3.6	100	.72	0.9	2.1	718
Garbage	7	7	−3.0	97	.75	0.2	2.7	739
Germs	4	5	−3.2	100	.74	0.5	2.3	710
Gift	44	4	3.4	100	.79	1.4	2.0	655
Gold	37	4	2.6	93	.65	1.8	1.4	766
Grease	12	6	−2.2	93	.63	0.9	1.6	889
Guns	142	4	−2.2	97	.58	1.2	1.9	719
Hangover	3	8	−2.9	93	.67	0.5	2.5	847
Hatred	20	6	−3.6	100	.73	0.8	2.3	783
Hawaii	16	6	3.4	97	.68	1.1	2.0	703
Hell	86	4	−3.5	93	.67	0.4	2.6	735
Hitler	16	6	−4.2	90	.78	0.1	2.8	730
Holiday	30	7	3.6	100	.75	0.8	2.4	651
Hornet	1	6	−2.4	93	.60	0.9	1.8	866
Ice cream	1	9	3.2	100	.75	1.8	1.3	620
Kitten	10	6	2.8	90	.68	0.9	1.8	731
Knives	7	6	−0.9	83	.50	2.2	0.9	894
Landlords	15	9	−0.6	83	.57	2.3	0.6	958
Litter	3	6	−3.1	93	.73	0.3	2.5	858
Liver	17	5	−1.4	77	.48	1.1	1.6	790

(*Appendix continues on next page*)

Appendix (*continued*)

Attitude object	Frequency[a]	Length[b]	Evaluation[c]	Consistency[d]	Consensus[e]	Ambivalence[f]	Polarization[g]	Latency[h]
Magazine	65	8	2.7	97	.77	1.7	1.4	787
Monday	72	6	−0.8	83	.46	1.3	1.5	890
Money	262	5	2.8	90	.64	2.4	1.1	780
Mosquito	2	8	−2.9	73	.70	0.2	2.2	763
Movies	60	6	3.2	97	.78	2.1	1.3	667
Music	216	5	4.2	97	.99	1.4	2.0	653
Parade	26	6	1.7	97	.67	1.6	1.3	724
Party	255	5	3.2	97	.73	1.9	1.3	685
Pie	19	3	1.7	87	.66	1.1	1.6	686
Pizza	3	5	2.8	100	.73	1.3	1.8	639
Priest	33	6	0.9	80	.48	1.4	1.5	929
Radiation	100	9	−3.2	90	.67	1.3	1.9	822
Rats	10	4	−3.7	97	.74	0.4	2.5	676
Rattlesnake	8	11	−2.7	93	.64	1.1	1.9	778
Reagan	62[i]	6	−0.3	73	.40	1.5	1.3	1,004
Recession	7	9	−1.7	80	.58	0.6	1.9	1,051
Russia	86	6	−1.0	83	.48	1.9	0.8	994
Silk	13	4	2.4	97	.72	0.6	1.9	749
Smoking	8	7	−3.1	97	.62	0.5	2.3	808
Snow	56	4	2.5	93	.66	2.1	1.2	818
Spider	2	6	−1.5	90	.55	1.4	1.3	882
Spinach	2	7	0.7	73	.47	1.1	1.5	855
Sports	48	6	2.3	97	.63	1.2	1.5	781
Stereo	13	6	3.7	97	.85	1.0	1.9	733
Storms	31	6	0.3	83	.47	2.0	0.8	878
Strawberries	2	12	2.9	93	.70	0.4	2.3	721
Summer	151	6	3.5	97	.76	1.1	1.9	663
Sunshine	8	8	3.9	97	.84	1.0	2.3	692
Swimming	37	8	2.6	93	.68	0.9	1.8	720
Taxes	43	5	−2.4	90	.63	1.8	1.3	763
Television	51	10	1.3	90	.58	2.2	0.8	761
Tequila	0	7	−0.5	87	.40	1.3	1.4	981
Toothache	1	9	−3.4	97	.78	0.3	2.5	804
Virus	13	5	−3.5	93	.78	0.5	2.5	746
Vodka	0	5	−0.6	77	.40	1.3	1.5	902
War	492	3	−4.1	90	.77	0.9	2.5	706
Weeds	5	5	−2.0	83	.65	0.6	2.0	828
Worms	8	5	−1.9	90	.58	1.4	1.3	766

[a] Frequency per million words from Francis and Kucera (1982) norms. [b] Number of letters. [c] Mean rating by sample of 274 New York University students, on scale from −5 (*extremely negative*) to 5 (*extremely positive*). [d] Percentage of subjects (n = 30) in normative consistency study who gave the same evaluation of the object in both experimental sessions. [e] Consensus factor score: mean of (a) the proportion of subjects across all experiments (n = 412) giving the modal evaluation of the object and (b) inverse of standard deviation of evaluation rating (n = 274). [f] Mean ambivalence score from normative study (n = 32); scores could range from 0 (low) to 6 (high). [g] Mean polarization score from normative study; scores could range from 0 (low) to 3 (high). [h] Mean latency (in milliseconds) to make the evaluative response for all subjects in the normative consistency study (first session), normative ambivalence study, and Experiments 1–3 (n = 168). [i] As the Francis and Kucera (1982) corpus is of material published in 1961, this number is the frequency of "Eisenhower" (there being more than one well-known "Kennedy" at the time).

Received May 18, 1990
Revision received June 11, 1991
Accepted November 25, 1991 ■