JOURNAL OF EXPERIMENTAL SOCIAL PSYCHOLOGY **32**, 104–128 (1996)
ARTICLE NO. 0005

The Automatic Evaluation Effect: Unconditional Automatic Attitude Activation with a Pronunciation Task

JOHN A. BARGH, SHELLY CHAIKEN, PAULA RAYMOND, AND CHARLES HYMES

New York University

Received: November 18, 1994; revised: April 22, 1995; accepted: May 4, 1995

Previous demonstrations of automatic attitude activation have used the same critical test of the presence of automaticity (Bargh, Chaiken, Govender, & Pratto, 1992; Chaiken & Bargh, 1993; Fazio, Sanbonmatsu, Powell, & Kardes, 1986). In this task, subjects first classify as quickly as they can each of a series of adjectives as to whether each adjective is good or bad in meaning. Subsequently, on each priming trial, a word corresponding to an attitude object briefly precedes the adjective target. Whether this attitude object *prime* automatically activates its corresponding evaluation in memory is indicated by whether the presence of that prime word on a trial facilitates or interferes with the conscious and intended evaluation of the adjective target. Although there is some disagreement as to the generality of the automaticity effect (Chaiken & Bargh, 1993; Fazio, 1993), all relevant experiments have obtained it for at least some of the subject's attitudes. The present experiments investigated the role that the adjective evaluation task itself may play in producing the automaticity effect. If the effect is truly unconditional, requiring only the mere presence of the attitude object in the environment, it should not depend on the subject consciously thinking in terms of evaluation. Three experiments designed to remove all such evaluative features from the basic paradigm revealed that the automaticity effect does not depend on such strategic evaluative processing. In fact, consistent with the *affect primacy* hypothesis (Murphy & Zajonc, 1993), removing conscious processing aspects from the paradigm serves to increase, rather than decrease, the generality of the automatic evaluation effect. © 1996 Academic Press, Inc.

The accessibility of information in memory—the ease with which it can be retrieved—is of critical importance to how one is able to relate

This research was supported in part by National Science Foundation Grant SBP-43095 and National Institute of Health Grant MH43265 to the first author. Experiments 1 and 2 were presented at the 1990 meetings of the Eastern Psychological Association, New York City. We thank Mark Zanna, Magda Garcia, Roger Gner-Sorolla, and Tony Greenwald for their comments on an earlier version. Address reprint requests to John A. Bargh, Department of Psychology, New York University, 6 Washington Place, Seventh Floor, New York, NY 10003.

0022-1031/96 $12.00

past experiences to present circumstances. Even if knowledge is available regarding with which hand a basketball opponent tends to shoot layups, unless that knowledge can be quickly accessed when (s)he is charging down the lane, it won't be of much use. Or if a friend suggests a certain restaurant for dinner, it would be nice to access right then the stored knowledge that you've never enjoyed a meal there, instead of realizing it later on when you've settled in and are pondering the menu. As these examples illustrate, we must frequently make choices and decisions quickly and "on the fly." Unless relevant information is readily accessible at that moment, its impact will be negligible.

The accessibility of knowledge is a function of how *recently* it has been used and also of how *frequently* it has been applied in the past (Higgins & Bargh, 1987; Higgins & King, 1981). Frequently used or *chronically* accessible knowledge exerts a greater influence on judgment than does other relevant but less accessible knowledge regardless of whether it has been recently used or not; it is in a relatively permanent state of increased activation (Bargh, 1989; Higgins, Bargh, & Lombardi, 1985; Higgins, King, & Mavin, 1982). The importance of chronic accessibility effects is their relatively unconditional nature (see Bargh, 1989)—they do not depend on the person having recently thought about or used the concept. In other words, chronically accessible mental representations become active upon the presence of relevant environmental information, even if intentional thought and attention are directed elsewhere (e.g., Bargh & Pratto, 1986).

Fazio (e.g., 1986, 1989) has applied accessibility logic to the question of which of one's attitudes will be most likely to influence behavior. Fazio defines attitude *strength* as the strength of the association in memory between the representation of the attitude object and one's evaluation of that object, and assumes that the stronger this association, the more accessible the attitude. According to the attitude accessibility model, the greater the accessibility of an attitude at the time of encountering the attitude object, the greater the likelihood that attitude will influence behavior toward the object at that time. Therefore, the strongest and most consistent attitude influences will come from those attitudes that are most chronically accessible, because such attitudes have fewer preconditions (e.g., recent thought about the attitude object) on their likelihood of activation. Fazio, Sanbonmatsu, Powell, and Kardes (1986) argued that a person's strongest or most chronically accessible attitudes will thus be capable of being activated automatically, upon the mere presence of the attitude object in the environment. Therefore, even if the person is not thinking about his or her feelings or opinion of the attitude object, that attitude will nonetheless become active and influence the person's behavior toward the object (for an alternative perspective on attitude–behavior consistency, see Aizen, 1996).

Fazio et al. (1986) adapted the sequential priming paradigm of Neely (1977) to test whether at least some attitudes become active in this automatic fashion.

In order to assess the chronic accessibility or strength of each subject's attitudes toward a variety of objects, subjects were asked to evaluate each object as quickly as they could in the first phase of the experiment. Fazio et al. (1986) operationalized attitude strength and, hence, chronic accessibility, in terms of how quickly subjects responded with their *good* or *bad* evaluation (by pressing the corresponding button on a response box); the faster the response, the stronger and therefore more accessible the attitude was presumed to be. The attitude objects corresponding to the fastest (i.e., strongest attitudes) and slowest (i.e., weakest attitudes) good and bad responses in this *attitude assessment* task were selected to serve as priming stimuli in the second task.

In the second phase of the experimental session, subjects were again instructed to evaluate as good or bad each of a series of stimuli. In this task, all of the target stimuli were adjectives, preselected to be clearly positive or negative in meaning. However, unlike the first attitude assessment task, on each trial of this adjective assessment task a prime word appeared for 250 ms just before the adjective target, too brief for any strategic or intentional activation of the attitude to occur (see Neely, 1977; Posner & Snyder, 1975). As noted above, these primes corresponded to each subject's strongest and weakest good and bad attitudes. Subjects were instructed to remember this first prime word on each trial (in order to repeat it out loud at the end of the trial), but their main instruction was to evaluate the second word (the adjective target) as quickly as possible.

If the attitude associated with the prime stimulus became active automatically upon the presentation of the prime, this *good* or *bad* evaluation would then be active during the time in which the subject classified the adjective target on that trial as good or bad. Although Fazio et al. (1986) did not provide an explicit mechanism for how the automatically activated evaluation would affect responses to the adjectives (an issue that becomes more important in light of the present findings; see General Discussion), Neely (1977) and Logan (1980) have discussed such effects in terms of response facilitation and competition. To the extent the subject's response to an adjective target (e.g., *beautiful*) is the same as the evaluation automatically triggered by the prime (e.g., *puppy*), it should take less time to evaluate the adjective because the correct response is already activated. To the extent the response automatically activated by the attitude object prime (e.g., *devil*) is the incorrect response to the adjective, it will take longer to evaluate the adjective due to the need to inhibit that competing response. Thus, to the extent the attitude object prime automatically activated its corresponding evaluation in memory, it would facilitate evaluating adjectives of congruent valence relative to adjectives of incongruent valence. If the corresponding attitude was not capable of automatic activation, the presence of that attitude object prime should have no effect on adjective evaluation latencies.

Fazio et al. (1986, pp. 231, 233, 235) predicted that subjects' strongest but not their weakest attitudes would become active automatically, because of their presumed difference in chronic accessibility. The results of two experiments (Experiments 1 and 2) supported this hypothesis: primes corresponding to the subject's strongest but not weakest attitudes reliably facilitated evaluation response times for adjectives of the same, relative to opposite, valence.

However, aspects of the Fazio et al. (1986) Experiment 3 results suggested that even weak attitudes may be capable of automatic activation. To Bargh et al. (1992), the reliable automatic activation found for the subjects' weak attitudes was a clue that the phenomenon was a relatively general one. They noted that a potentially critical difference between the Fazio et al. (1986) Experiments 1 and 2, which did not report automaticity in the weak attitude condition, and their Experiment 3 which did, was that subjects in the former experiments had just been asked to think about and give their attitudes toward each of the attitude objects. Because this recent thought about one's attitudes was a theory-irrelevant feature of the paradigm—which was intended to test whether attitudes become active upon the *mere presence* of the attitude object—Bargh et al. (1992, Experiment 2) placed a 2-day delay between the attitude assessment and adjective evaluation tasks. Under these conditions, the automatic attitude activation effect occurred for the subjects' weakest as well as strongest attitudes. In fact, the size of the effect was equivalent in the two cases. Chaiken and Bargh (1993) replicated this result in an experiment which varied the time delay between the attitude assessment and adjective evaluation tasks. Thus, making the automaticity paradigm more similar to "mere presence" conditions in which a person has not recently thought about one's feelings towards the object, serves to increase the ubiquity of the automaticity effect and, apparently, to remove the moderating influence of attitude strength.

The primary purpose of the present research is to continue our examination of the *unconditionality* of automatic attitude activation. Bargh et al. (1992) noted several aspects of the original paradigm that might have contributed artifactually to producing the effect, such as instructions to hold the prime word in memory as well as having subjects evaluate each of the attitude objects prior to the test of automaticity. As other studies had shown that temporary accessibility (i.e., priming) effects are qualitatively identical to chronic accessibility effects (Bargh et al., 1986, 1988), it was important to rule out a causal role for the attitude assessment task. As noted above, the Bargh et al. (1992) Experiment 2 and its replication by Chaiken and Bargh (1993) demonstrated the automaticity effect even when the assessment task occurred two days before the test of automaticity.

Before concluding from these results that automatic attitude activation is truly unconditional, another potentially critical feature of the paradigm needs to be addressed. In the original Fazio et al. (1986) paradigm, not only did subjects evaluate each of the attitude object stimuli prior to the test of attitude automaticity but also the test of automaticity itself involved a conscious evaluation of adjective targets. Thus, although Bargh et al. (1992) ruled out the initial attitude assessment task as a cause of the automaticity effect (if anything, the effect was more general without it), their three experiments and the Chaiken and Bargh (1993) replication also required subjects to consciously evaluate adjectives in the task assessing the automaticity of their attitudes toward the priming stimuli.

In other words, in the task that assessed whether attitudes become active upon the mere presence of the attitude object, subjects had the conscious, intentional goal of evaluation. They had their fingers on buttons labeled "good" and "bad" and were trying to evaluate each of a series of adjectives as quickly as possible. It would be surprising, given this intensive evaluative processing goal or mindset (see Gollwitzer, 1990; Smith & Branscombe, 1987; Smith & Lerner, 1986), if subjects did not *also* evaluate the prime word that appeared for a quarter of a second before the adjective target. Not to do so would require subjects to "turn off" their evaluative processing goal for that short period of time and then turn it on again immediately when the adjective appeared. As Neely (1977) showed that it takes at least 500 ms for conscious processing strategies to become fully active, it seems dubious that subjects engaged the evaluative goal only upon presentation of the adjective target on each trial. More probably, they were thinking in terms of goodness and badness during presentation of the attitude object prime as well as the target.

The importance of whether attitudes become activated on the mere presence of the attitude object is that such automatically activated attitudes will exert a more consistent influence over behavior (Fazio, 1986, 1989). "Mere presence" means just that—no need for the person to have the intention or goal to evaluate the object or to have to think about how he or she feels about it. In the automaticity task used in the Fazio et al. (1986), Bargh et al. (1992), and Chaiken and Bargh (1993) studies, subjects did not intend to evaluate the object itself, but they nonetheless had a conscious evaluation goal. The question then becomes whether the evidence thus far of automatic attitude activation is dependent on the subject having just such a conscious evaluative mindset.

In the present Experiment 1, we removed the adjective evaluation task from the paradigm so that the hypothesis of automatic attitude activation in the absence of a strategic evaluation processing goal could be tested. Instead of evaluating each of a series of target adjectives as good or bad, subjects merely pronounced them as quickly as they could. Balota and Lorch (1986) have shown that this pronunciation task is a very sensitive instrument in assessing spreading activation effects in memory. In particular, they found that in the sequential priming paradigm developed by Neely (1977) and used in the attitude automaticity research, the pronunciation task was capable of detecting two steps of activation spread compared to just one for the lexical decision task. For example, in the lexical decision task *lion* primed *tiger* and *tiger* primed *stripes*, but *lion* did not prime *stripes*; in the pronunciation task, this two-step *lion–stripes* priming effect was detected.

In all but one other respect, Experiment 1 was a replication of the Fazio et al. (1986) Experiment 2 and Bargh et al. (1992) Experiment 1; as Bargh et al. (1992) found that removing the instructions to the subject to remember the priming stimuli did not affect the results (as well as being theory-irrelevant), those instructions were deleted from the present experiments as well.

EXPERIMENT 1

Method

Subjects. Forty-seven students enrolled in the Introductory Psychology course at New York University (NYU) participated in the experiment in partial fulfillment of a course requirement. Data from 4 of these students were excluded from the analyses because they did not meet our criterion for English language ability, which was learning to speak English by the age of 5.

Materials and apparatus. Subjects were shown into the 2.7 $\times$ 3 m experimental room and seated in front of a cathode-ray-tube (CRT) display, which was under program control of an Apple II Plus microcomputer. In front of the subject was a two-button response box that was connected as an input device to the computer during the initial attitude-assessment task of the experiment. Subjects were instructed to rest their hands on the box so that their left-hand index finger was poised above the left (labeled *bad*) button, and their right-hand index finger was positioned above the right (labeled *good*) button.

For the pronunciation task that followed, a microphone connected to a Scientific Prototype voice-activated relay served as the input device to the computer. The sensitivity and gain controls of the voice-activated relay were preset for all subjects so that when the subject spoke a target word out loud into the microphone, a signal would be sent by the relay to the computer's input port.

All other materials and equipment were the same as in the Fazio et al. (1986) Experiment 2 and Bargh et al. (1992) Experiment 1. For the initial attitude assessment phase of the experiment, each subject was randomly assigned to one of four random presentation orders of the 92 attitude objects used in the Fazio et al. (1986) and Bargh et al. (1992) research (order did not interact significantly with any of the results reported below and will not be discussed further). For the subsequent priming task, the 10 positive and 10 negative adjectives used in that previous research again served as the target stimuli. Each of these target adjectives was paired with one of each of the five types of priming stimuli (fast–good, fast–bad, slow–good, slow–bad, and baseline three-letter strings), creating a total of 100 experimental trials. Thus, each target adjective was pronounced by subjects five times during the experimental session, so that its relative speed of pronunciation when preceded by each of the five prime types could be assessed.

A random ordering of these 100 prime-target combinations was created, but with the following restrictions: each target adjective appeared once in each block of 20 trials, each of the 20 primes appeared in a different random order in each block, and within each block, 2 of each prime type were paired with positive targets and 2 with negative targets (see Bargh et al., 1992, p. 899).

Procedure. Subjects were told that the experiment concerned judgments about words, and that they would be performing two different tasks. In the first *attitude-assessment* task, they were instructed to evaluate each of a series of words presented on the computer screen using the buttons labeled *good* and *bad*. Subjects were told to be sure to make the correct response but also to make it as quickly as possible. Prior to the start of the experimental trials, subjects were given 10 practice trials with the experimenter present. Then the experimenter left the room, and the computer presented the subject with one of the four random orderings of the 92 attitude object stimuli. Each of the stimulus words was presented in the middle of the screen and remained there 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, and after a 3-s pause, presented the next stimulus word.

After completing the attitude assessment task, subjects were given a 3-min rest period during which the computer program sorted their responses in ascending order of latency, separately for the *good* and *bad* responses. Next, the program selected the attitude object names corresponding to the subject's four fastest and four slowest *good* responses, and four fastest and four slowest *bad* responses, and entered these words into the list of prime stimuli in the manner described above in the Materials and Apparatus section.

Also during this time, the experimenter removed the response button box and replaced it with the microphone. Subjects were told that in the second part of the experiment they would be pronouncing words as quickly as they could into the microphone in front of them. They were further informed that unlike the first task, this time there would be two words presented on each trial, one right after the other, and that they were to pronounce the second of these words as quickly as they could after it was presented. Again, accuracy as well as speed was emphasized. Subjects were given 10 practice trials to ensure they understood the procedure before the experimenter left the room and started the experimental trials.

As in the Fazio et al. (1986) and Bargh et al. (1992) experiments, the prime word appeared in the middle of the CRT screen for 200 ms, with a 100-ms blank-screen interval, followed by presentation of the target adjective in the same location. The target word remained on the screen until the subject began to pronounce it. The computer program recorded the latency of the pronunciation response, and the experimenter monitored its accuracy from an adjacent room (there were no incorrect responses in the experiment). There was a 4-s pause before the start of the next trial. When all 100 trials had been completed, the experimenter debriefed the subject and thanked him or her for participating.

Results

The 43 subjects made a total of 4300 responses in the pronunciation task. Following the procedure of Balota and Lorch (1986; see also Ratcliff, 1993) to reduce the effect of outlier latencies on the analyses, those under 300 ms (0.2%) or over 1000 ms (0.3%) were excluded. In nearly all cases, the shortest latencies corresponded to trials on which the subject made an extraneous sound (e.g., a cough) that triggered the relay, and the longest latencies to trials on which the subject spoke too softly for the relay to detect the response (forcing the subject to repeat the stimulus word in order to erase it from the screen and proceed to the next trial).

The mean pronunciation latency (after exclusion of outliers) was 518 ms ($SD = 92$), and the median latency was 501 ms. Because of the positive skewness of the pronunciation latency distribution, the raw latencies were transformed in two different ways in order to eliminate the skew problem prior to computing analyses of variance (ANOVAs). The logarithmic transform is recommended by Winer (1971, p. 400) for reducing positive skew in latency data, whereas the reciprocal transform is favored by Fazio (1993). (See also the comparisons of these and other transformation methods by Box, Hunter, & Hunter, 1978, pp. 231–241, and Ratcliff, 1993, pp. 510–532). The results reported below are from the analysis of the logarithmically transformed data, but virtually identical results were obtained in the analysis of the reciprocal data, and the same pattern of significant and nonsignificant results was obtained in both analyses (see Footnotes 2, 4, and 5).

We computed each subject's mean (log-transformed) response latency for each of the 8 cells of the design (omitting the baseline prime condition; see Footnote 1), formed by the complete crossing of the within-subjects factors of Prime Valence (positive vs negative), Target Valence (positive vs negative) and Attitude Speed (fast vs slow attitude assessment response). The mean (untransformed) latencies across subjects are shown in Fig. 1.[1] The 2 × 2 × 2 repeated-measures ANOVA on these scores revealed a highly significant Prime Valence × Target Valence interaction, $F(1, 39) = 44.51$, $p < .0001$ ($MS_E = .0012$). The presence of this significant interaction and the pattern of means shown in Fig. 1, with shorter latencies for the conditions in which prime and target valences are congruent (shaded bars) and relatively longer latencies for the

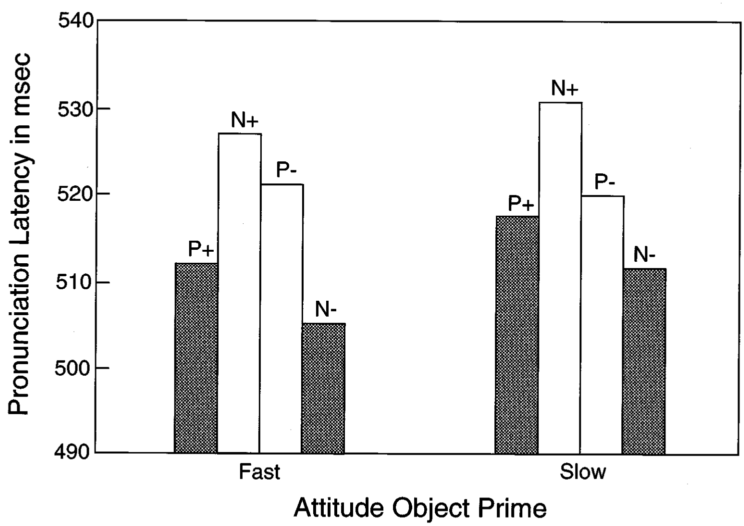

FIG. 1 Mean (untransformed) target pronunciation latencies (in milliseconds) by prime type, prime valence (P, positive; N, negative) and target valence (+, positive; −, negative), Experiment 1. Shaded bars represent conditions in which prime and target valences were congruent.

[1] Presentation of the mean response latencies rather than a facilitation score in which the mean is first subtracted from a baseline or control prime condition mean is a departure from the Fazio et al. (1986) and Bargh et al. (1992) articles. However, because the same baseline mean was used as a constant within each target valence condition (i.e., there was a baseline/positive target condition and a baseline/negative target condition), the pattern of means and the sizes of the differences between the means are identical whether raw means or facilitation scores are presented. We chose to dispense with the facilitation score method because of difficulties of creating a true baseline or "neutral" prime condition (see discussion in Fazio et al., 1986; Bargh et al., 1992) and the consequent lack of justification for presenting absolute facilitation and inhibition scores. Without some certainty as to an appropriate baseline measure, the safest course is to interpret the means as faster or slower relative to each other.

valence-incongruent conditions within each target valence condition, is the signature of the automatic attitude activation effect.

As is evident from Fig. 1, the automaticity effect held equally strongly for subjects' fastest or strongest attitudes and for their slowest or weakest attitudes. The critical test of whether the automaticity effect varies as a function of attitude strength, at least as indicated by attitude response latency, is the Attitude Speed × Prime Valence × Target Valence interaction, which was not significant, $F(1, 39) = 0.76$.[2]

Of lesser theoretical importance were main effects for Attitude Speed, $F(1, 39) = 5.10$, $p = .03$ ($MS_E = .001$), and Target Valence, $F(1, 39) = 19.31$, $p < .0001$, ($MS_E = .001$). Subjects were faster to pronounce words when they were preceded by "fast" (vs "slow") attitude object primes, and were faster to pronounce negative (vs positive) target adjectives. All other effects were nonsignificant (ps > .23).

Discussion

These results indicate that the automatic attitude activation effect is not conditional on the subject having an explicit, conscious evaluative goal. Removing the evaluative goal from the paradigm did not eliminate the automaticity effect. Furthermore, as in the studies by Bargh et al. (1992) and Chaiken and Bargh (1993), the elimination of this theory-irrelevant feature of the paradigm resulted in an automaticity effect that was equally probable for the subject's idiosyncratically strong and weak attitudes.

Although it is certainly true that one can never prove the null hypothesis (but see General Discussion), the absence of any hint of moderation of the automaticity effect by attitude strength, as operationalized in terms of attitude assessment latency, is telling. We have now failed to obtain evidence of such moderation in three experiments that removed theory-irrelevant features from the paradigm in order to more closely approximate "mere-presence" conditions. In the earlier studies, separating the attitude assessment and adjective evaluation tasks in time was sufficient to eliminate statistically reliable moderation by attitude strength (Bargh et al., 1992, Experiment 2; Chaiken & Bargh, 1993). In the present study, eliminating the conscious evaluation goal during the automaticity test was by itself sufficient to eliminate moderation. The only conditions under which we have obtained significant moderation of the automaticity effect by attitude strength were those in which we followed the original Fazio et al. (1986) procedure exactly (Bargh et al., 1992, Experiments 1 and 3; Chaiken & Bargh, 1993, immediate assessment condition). Moderation of the effect by attitude strength thus appears to occur mainly under procedural conditions that involve recent or concurrent strategic evaluative processing—that is, more than just the mere presence of the attitude object.

[2] The results of the ANOVA on the reciprocally transformed latencies revealed the same reliable Prime Valence × Target Valence interaction, $F(1, 39) = 45.25$, $p < .0001$ ($MS_E = 41.02$), and nonsignificant Attitude Speed × Prime Valence × Target Valence interaction, $F(1, 39) = 0.05$.

EXPERIMENT 2

Although we eliminated explicit instructions to evaluate the target adjectives from the paradigm in Experiment 1, evaluative aspects of the experimental session remain. As in the original paradigm, subjects evaluated each of the 92 attitude objects as quickly as they could immediately prior to the adjective pronunciation task. Thus subjects were thinking for several minutes about their likes and dislikes, just prior to pronouncing the adjectives and being exposed to some of those same attitude objects as priming stimuli.

There are two ways in which this immediately prior attitude assessment task might have contributed to artifactually producing evidence of automaticity. The first, that subjects' attitudes might be temporarily activated by the assessment task, has already been ruled out as a causal influence by the Bargh et al. (1992, Experiment 2) and Chaiken and Bargh (1993) delay data. But the *evaluative nature* of the attitude assessment task could itself have carried over to the adjective pronunciation task, either through extensive priming of the concepts of *good* and *bad* (thereby causing the attitude object primes to be categorized using those concepts; e.g., Bargh, 1989; Higgins, 1989), or by a priming and consequent residual influence of the evaluative goal or mindset itself.

Gollwitzer, Heckhausen, and Steller (1990) demonstrated how a processing goal or mode used in one context can carry over to subsequent contexts and influence processing even though there is no explicit choice of that goal in the second context. Their subjects were asked to adopt either a "deliberative" or an "implemental" mindset in a first experiment, by considering a personal problem either by thinking about what the best way to solve it would be, or by making specific plans as to what they were actually going to do about it. In the second, ostensibly unrelated experiment, subjects completed a fairy tale after being given only the first few sentences (e.g., a king had to go to war but did not want to leave his daughter unprotected). Subjects who had previously engaged in deliberative thought were more likely to complete the story using deliberations (e.g., all the possibilities the king was thinking about), while subjects who had previously an implemental mindset were more likely to complete the story with actional content (e.g., what the king did to solve the problem). Thus, the activated processing mode carried over to a subsequent task and to informational input for which it was not intended.

In Experiment 2, we sought to eliminate the possible contamination of the pronunciation test of attitude automaticity by removing the attitude assessment task from the experimental design. To do so, we first had to preselect the strong and weak, good and bad attitude object primes for all subjects. We did this by consulting the Bargh et al. (1992, Appendix) normative data for each of the 92 attitude object stimuli, and selecting attitude objects that possessed characteristics associated with attitude strength (see Bargh et al., 1992). Normatively *strong* attitude object primes were those that were characterized by relatively fast response latencies, extreme evaluations, low ambivalence ratings, and high consistency of evaluation, whereas normatively *weak* attitude object primes were characterized by slow response latencies, moderate evaluations, high ambivalence ratings, and inconsistency of evaluation. In selecting these sets of prime stimuli, we made every attempt to select the strongest and the weakest possible attitude objects in order to assess the automaticity effect across the entire spectrum of attitude strength.

Method

Subjects. Thirty-one NYU undergraduates enrolled in the Introductory Psychology course participated in partial fulfillment of a course requirement. Data from 6 subjects were excluded from the analyses because they did not meet the English-language criterion, leaving a final sample of 25 participants.

Materials and apparatus. Other than the fact that the good-bad response box was not present, and that the microphone was present from the beginning of the experimental session, the experimental room and apparatus were the same as for Experiment 1. The same random order of prime type/target word combinations was used as before. However, the particular attitude object names that served as the fast–good, fast–bad, slow–good, and slow–bad primes were predetermined based on the normative data of Bargh et al. (1992), instead of being determined idiosyncratically by each subject in an attitude assessment phase.

Separate sets of "strong–good," "strong–bad," "weak–good," and "weak–bad" attitude object primes were preselected based on their normative *evaluative latency*, *evaluative extremity*, *ambivalence*, and *consistency of evaluation*—variables all associated with attitude strength (Eagly & Chaiken, 1993; Krosnick, Boninger, Chuang, Berent, & Carnot, 1993) and all of which were significantly intercorrelated in the Bargh et al. (1992) normative study. We selected as *strong* good and bad primes those that were among the most extremely evaluated, associated with the least ambivalent feelings, and the most quickly and consistently evaluated of the 92 attitude object stimuli. The *weak* good and bad primes were among the least extremely evaluated, associated with the most ambivalent feelings, and the most slowly and least consistently evaluated of the attitude objects.[3] Table 1 presents these preselected sets of primes.

These primes were placed in the same trial order positions reserved for the fast–good, fast–bad, slow–good, and slow–bad primes in Experiment 1 so that, as before, each target adjective was paired with one of each type of priming stimuli, giving a total of 80 experimental trials. We created two prime stimuli orderings so that which particular member of a prime-type set appeared on a given trial varied between the two lists, and subjects were randomly assigned to one of these order conditions. (Again, order made no difference to the reported results.)

Results

As in Experiment 1, latencies less than 300 ms (1.1%) or greater than 1000 ms (2.4%) were omitted from the analyses. (Although the percentage of

[3] Contrary to Krosnick et al.'s (1993) conclusion that these and other indicators of attitude strength should be viewed as distinct dimensions of this construct, Pomerantz, Chaiken, and Tordesillas (1995) recently found evidence to support at least two general aspects of attitude strength. More important for present purposes, their factor analytic findings support the validity of our aggregated, normative index of attitude strength: attitude extremity, response latency, and ambivalence all had high loadings on the factor that Pomerantz et al. (1995) labeled Commitment, and low loadings on the factor they labeled Embeddedness (consistency of evaluation was not examined in that research).

TABLE 1
PRESELECTED STRONG AND WEAK ATTITUDE OBJECT PRIMES (EXPERIMENTS 2 AND 3)

Prime type	Evaluation	Consistency	Ambivalence	Latency
Strong–Good				
Friend	4.4	100	1.0	669
Dancing	3.1	100	0.3	661
Flowers	3.5	100	0.6	671
Holiday	3.6	100	0.8	651
Strong–Bad				
Cancer	– 4.4	97	1.8	681
Death	– 3.1	97	0.9	724
Funeral	– 3.6	100	0.9	718
Disease	– 3.9	100	0.3	708
Weak–Good				
Dormitory	1.1	73	2.1	1014
Spinach	0.7	73	1.1	855
Priest	0.9	80	1.4	929
Pie	1.7	87	1.1	686
Weak–Bad				
Dentist	– 1.1	77	2.3	892
Exams	– 1.2	77	2.1	896
Liver	– 1.4	77	1.1	790
Vodka	– 0.6	77	1.3	902

Note. Evaluation extremity means for the 92 attitude objects (on a -5 to $+5$ scale) ranged from 0.3 to 4.4 for the positively evaluated objects, and from -0.3 to -4.4 for the negatively evaluated objects. Consistency refers to the percentage of subjects in the Bargh et al. (1992) normative study who gave the same evaluation of the object on two different occasions, with means ranging from 73 to 100. Ambivalence rating means ranged from 0.0 to 2.4 (on a 0–4 scale), and latency means ranged from 620 to 1051.

latencies greater than 1 s was higher than that for Experiment 1, it is comparable to the rate of 1.9% obtained in the Balota & Lorch, 1986, pronunciation experiment). Notwithstanding this trimming procedure the latency distribution was again positively skewed, with a mean of 522 ms ($SD = 102$) and median of 490 ms. As in Experiment 1, we performed a logarithmic transformation on the raw latencies before submitting them to the ANOVA.

The test of automatic attitude activation, the Prime Valence × Target Valence interaction, was again highly reliable, $F(1, 23) = 25.98$, $p < .0001$ ($MS_E = .003$). It can be seen in Fig. 2 that the signature pattern of faster pronunciation times when targets are preceded by primes of like valence than preceded by primes of opposite valence was obtained for both the strong and the weak prime condition. As in Experiment 1, the size of the automaticity effect did not vary as a function of attitude strength; the Attitude

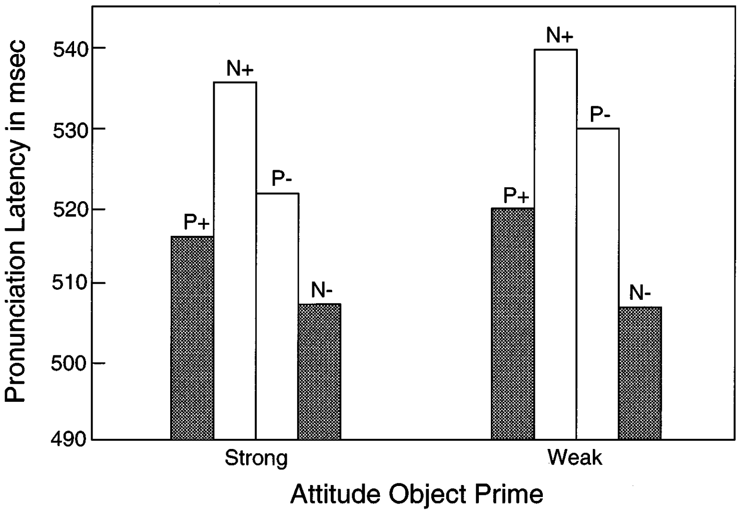

FIG. 2 Mean (untransformed) target pronunciation latencies (in milliseconds) by prime type, prime valence (P, positive; N, negative) and target valence (+, positive; −, negative), Experiment 2. Shaded bars represent conditions in which prime and target valences were congruent.

Strength × Prime Valence × Target Valence interaction was not significant, $F(1, 23) = 0.73$.[4]

Finally, as in Experiment 1, the main effect for Target Valence was reliable, $F(1, 23) = 13.49$, $p = .001$ ($MS_E = .002$), with subjects pronouncing the negative target adjectives more quickly than the positive adjectives. All other effects were nonsignificant at $p > .20$.

Discussion

The results of Experiment 2 replicate the essential findings of the first experiment, that removing aspects of the paradigm that might induce an evaluative processing strategy—and thereby produce misleading evidence of automatic evaluation—does not remove the automatic attitude activation effect. The case for unconditional automatic evaluation effects, in which environmental stimuli are classified as good or as bad immediately, efficiently, and uncontrollably by the individual, is strengthened by these results. Once again, making the paradigm more closely resemble the unconditional mere presence of the attitude object, in the present experiment by not requiring the

[4] The corresponding findings from the ANOVA on the reciprocally transformed latencies were virtually the same: for the Prime Valence × Target Valence interaction, $F(1, 23) = 27.07$, $p < .0001$ ($MS_E = 35.88$); for the Attitude Strength × Prime Valence × Target Valence interaction, $F(1, 23) = 0.92$.

subject to have recently evaluated or otherwise thought about the relevant attitude objects, continues to result in an automatic activation effect of substantial size.

As in Experiment 1, this automaticity effect was not moderated by attitude strength. The lack of evidence for moderation that we obtained under such nonevaluative conditions is all the more dramatic because we had preselected the "weak" attitude objects to be the weakest of the set of 92 attitude objects: the slowest evaluation latencies, the least consistent, the most ambivalent, and the least extremely evaluated. Yet these very weak attitudes were just as capable of automatic activation as the very strongest of the set of 92 objects, with no indication that the probability of automatic activation was greater for those so much higher in the several attitude strength qualities.

EXPERIMENT 3

We can identify one final aspect of the basic paradigm that may induce an evaluative mindset—the strongly valenced nature of the adjectives that subjects respond to during the test of attitude automaticity. It is not unreasonable to suppose that repeatedly seeing and pronouncing adjectives such as *beautiful*, *phony*, *horrible*, and *wonderful* could passively prime the concepts of good and bad, or an evaluative processing goal. It is also possible that subjects consciously notice the valenced nature of the target stimuli and infer that the experiment has something to do with evaluation.

To be completely confident that subjects automatically evaluate stimuli without any explicit or implicit cues to engage in evaluative processing, we replaced the strongly valenced adjectives with targets of less obvious valence. These target words were selected to be mildly positive and negative in nature, according to likability norms collected by Bellezza, Greenwald, and Banaji (1986). The moderate quality of the targets' evaluations would make it very unlikely that they would induce an evaluative processing goal.

Method

Subjects. Thirty-seven NYU Introductory Psychology students participated in the experiment in partial fulfillment of a course requirement. Data from 7 participants who did not meet our criterion for English language ability were excluded from the analyses.

Materials and apparatus. The only change from Experiment 2 was the substitution of a different set of target stimuli for the adjectives previously employed in the Fazio et al. (1986), Bargh et al. (1992), Chaiken and Bargh (1993), and present Experiments 1 and 2. Bellezza et al. (1986) had 76 subjects rate each of 399 words as to their pleasantness on a 1 (very unpleasant) to 5 (very pleasant) scale. Subjects also rated each word for its frequency of usage (in the subject's experience), and its ability to elicit a clear and vivid image.

As our purpose in Experiment 3 was to eliminate as much as possible any explicit or implicit evaluative quality from the experimental paradigm, we selected target stimuli that were not extremely or even noticeably positive or negative in valence. ·For the negative stimuli we chose nouns with normative pleasantness ratings around the scale value of 2, one unit away from the scale midpoint. Similarly, the positive target stimuli were selected to have scale values around 4.

TABLE 2
TARGET STIMULI FOR EXPERIMENT 3

	Pleasantness[a]	Frequency[b]
	Positive targets	
Agility	3.99	2.29
Bird	4.03	3.56
Brother	4.22	3.89
Dress	3.95	3.78
Earth	4.11	3.38
Eat	4.12	4.71
Exercise	4.01	4.13
Fur	4.09	2.29
Knowledge	4.11	3.71
Lake	4.21	3.24
	Negative targets	
Afraid	2.05	3.78
Army	2.35	3.10
Debt	1.75	3.03
Fire	1.87	3.29
Pest	1.85	2.62
Tobacco	2.14	3.21
Slap	1.79	2.55
Trouble	2.13	3.81
Useless	1.85	3.19
Weapon	2.08	3.11

[a] Mean pleasantness ratings by subjects in Bellezza et al. (1986) normative study (1 = very unpleasant, 5 = very pleasant).

[b] Mean frequency of word use ratings from Bellezza et al. study (1 = very infrequently, 5 = very frequently).

We also attempted to equate the positive and negative target stimuli as to their average frequency of use by the subjects. The resultant sets of positive and negative targets are shown in Table 2.

Within the set of positive targets and separately within the set of negative targets, each target word was randomly assigned an order number that specified that target's presentation positions in the sequence of 80 experimental trials. The same randomized target word presentation order from Experiments 1 and 2 was used. Thus, each target word appeared four times, preceded once each by each of the four types of prime stimuli (i.e., strong–good, strong–bad, weak–good, weak–bad). The same sets of primes were employed as in Experiment 2, and subjects were randomly chosen to receive one of the two prime presentation orders. As in the previous experiments, this order factor did not significantly interact with any of the reliable effects reported below.

Procedure. The procedure was identical to that of Experiment 2. Subjects were informed that they were participating in a language experiment, and were instructed to pronounce each of a series of words as quickly as they could. They were told that on each trial they would see two words, one right after the other, and that their task was to pronounce the second word as quickly as possible. Before engaging in the experimental trials a series of 10 practice trials was presented; both the primes and targets for these practice trials were taken from the Bellezza et al. (1986) norms following the same criteria as for the experimental target stimuli (e.g., *noise*, *small*).

Results

Again, pronunciation latencies less than 300 ms (0.9%) or greater than 1000 ms (1.6%) were omitted from the analyses. The grand mean of all pronunciation latencies in the experiment was 500 ms ($SD = 100$), and the median latency was 479 ms. Just as in Experiments 1 and 2, the trimmed latency distribution was characterized by significant positive skew, and so we performed a logarithmic transformation on the raw latencies before submitting them to the ANOVA.

Of central importance is the Prime Valence × Target Valence interaction. Once again it proved highly reliable, $F(1, 28) = 21.64$, $p < .0001$ ($MS_E = .001$). Figure 3 presents the mean response latencies, separately for the strong and the weak prime conditions. The signature pattern of the automatic attitude activation effect was again obtained, as subjects were faster to pronounce the nonadjective targets when they were preceded by primes of the same valence than when those targets followed primes of opposite valence. It is clear from Fig. 3 that the size of the automaticity effect did not vary with attitude strength, an impression confirmed by a nonsignificant Attitude Strength × Prime Valence × Target Valence interaction, $F(1, 28) = 0.01$.[5]

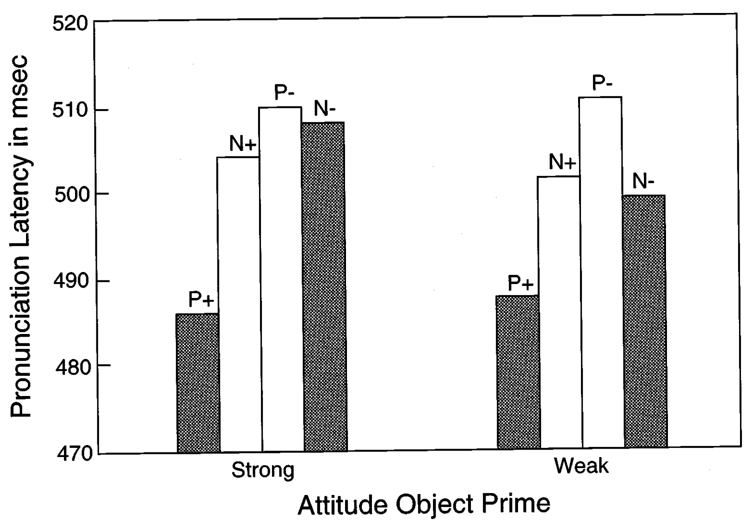

FIG. 3 Mean (untransformed) target pronunciation latencies (in milliseconds) by prime type, prime valence (P, positive; N, negative) and target valence (+, positive; −, negative), Experiment 3. Shaded bars represent conditions in which prime and target valences were congruent.

[5] In the ANOVA of the reciprocally transformed latencies, the Prime Valence × Target Valence interaction was reliable, $F(1, 28) = 17.94$, $p < .0002$, whereas the Attitude Strength × Prime Valence × Target Valence interaction was not reliable, $F(1, 28) = 0.21$.

As in the two previous experiments, there was a reliable main effect for Target Valence, $F(1,28) = 14.13$, $p = .0008$ ($MS_E = .003$). However, in contrast to the previous experiments, here the positive targets were responded to more quickly than the negative targets. All other effects were nonsignificant (*p*s > .12).

Discussion

We can now be quite confident that the automatic attitude effect is not conditional on subjects' having a concurrent evaluative processing goal, or in any way thinking more than usual in evaluative terms. When every possible evaluation-inducing aspect of the paradigm is removed, the automaticity effect continues to obtain. Moreover, once again the effect showed no sign of being moderated by attitude strength, replicating the findings of the present Experiments 1 and 2 as well as Bargh et al. (1992, Experiment 2) and Chaiken and Bargh (1993) in which the paradigm was modified to make it more closely approximate naturalistic conditions of mere presence of the attitude object.

GENERAL DISCUSSION

In three experiments, the successive removal of all evaluative aspects of the automatic attitude activation paradigm (Fazio et al., 1986) resulted in evidence of a universal and unconditional automatic evaluation response. All attitude object stimuli studied were shown to trigger an immediate, reflexive, and uncontrollable good or bad response, depending on the subject's evaluation of them. Subjects participated in what they thought was a study on language, and were instructed merely to pronounce each of a series of target words as quickly as they could. Under such minimal conditions, weak attitudes (e.g., toward *priest* or *vodka*) and strong attitudes (e.g., toward *friend* or *cancer*) alike were shown to become active automatically.

Is the Attitude Automaticity Effect Moderated by Attitude Strength?

The present results cohere with and extend our previous findings (Bargh et al., 1992, Experiments 2 and 3; Chaiken & Bargh, 1993) in which the mere presence of the attitude object resulted in its automatic evaluation, regardless of the strength of the attitude. Without the strategic evaluative features of the original paradigm, the automaticity effect is not reliably moderated by attitude strength. In fact, the *F* values for the test of the moderated automaticity effect were trivially small in each of the present experiments. Moreover, effect size calculations (using *g* and then converting to *d*, to correct for small sample bias; Hedges & Olkin, 1985) indicated that, across the three experiments, there was a negligible trend for the automaticity effect to be smaller for fast/strong versus slow/weak attitude object primes ($d = -.0073$, $p = .95$). Indeed, only Experiment 2 yielded a pattern of means and an effect size directionally consistent with the moderated automaticity effect ($d = +.0585$, $p = .84$; for Experiments 1 and 3, *d*s = $-.0462$ and $-.0063$; *p*s > .83).

Given the general equivalence of chronic accessibility and probability of automatic activation (Bargh, 1984), it would seem problematic that the attitude automaticity effect does not reliably covary with attitude accessibility. However, there are several reasons to believe that the automatic evaluation effect is "special" and does not obey the same rules as non-evaluative semantic processing.

Evaluative versus semantic priming effects: Why does **priest** *prime* **fur**? Our first point is that the present findings of evaluative priming cannot be explained in terms of existing models of spreading semantic association. In the earlier studies of Fazio et al. (1986) and Bargh et al. (1992), which used adjective evaluation latencies as the dependent measure in the test of automaticity, it was quite plausible that the effect of prime on target was one of response competition or inhibition or both. That is, because the response to the adjective target was either "good" or "bad," automatic activation of the concept "good" or "bad" by the attitude object prime would either facilitate or inhibit this response (Logan, 1980; Neely, 1977; Posner & Snyder, 1975). If the correct response to the adjective was "good," for example, and the attitude object prime had already automatically activated or primed the response "good," then less controlled and attentional processing would be needed to make the correct response; on the trials on which prime and target valence mismatched, the response automatically activated by the prime would be incorrect and would require attentional resources and time to be inhibited. This is the same principle invoked in the Stroop color-word paradigm (Bargh & Pratto, 1986; Logan, 1980).

However, in the pronunciation task the correct response is neither "good" nor "bad"; it is the target word itself. How then can we explain why brief exposure to the word "priest" rather than "liver" makes subjects say "fur" faster? How does activation get from the concepts *priest* to *fur* so quickly?

In some manner, to produce the present results it must be the case that a positively (negatively) evaluated object causes all other positively (negatively) evaluated object representations to become active in memory. This must be occurring because the primes and targets in the present experiments were randomly paired, and they shared no semantic feature other than valence. Any theoretical account of the automatic evaluation effect must be able to explain how the prime *flowers* causes one to be able to say the word *knowledge* faster (both sharing a positive evaluation), and *dentist* makes one faster to say *army* (both sharing a negative evaluation).

The pervasiveness of the evaluative priming effect cannot be easily explained in terms of spreading activation along a semantic associative network. According to spreading activation models of semantic memory (e.g., Collins & Loftus, 1975; Lorch, 1982; Neely, 1977; Posner & Snyder, 1975), the probability of activation spreading from one concept to another is a function of the number of features held in common by the two representations, relative to the number of other representations that contain those features. Posner and

Snyder (1975), for example, argued that a stimulus automatically activates its associated representation in memory, with this activation then spreading to other representations "nearby" in the feature space but not to more remote, semantically unrelated representations. Thus, *automobile* and *truck* will prime each other because they share many features that most other representations do not contain. However, *apple* and *fire* should not prime each other, because although both representations contain "red" as a central feature, so do many other representations; moreover "red" is the only feature held in common. In the automatic attitude paradigm, primes and targets share (or do not share) valence and no other feature. Given that most representations probably contain the feature "good" or "bad," the phenomenon of priming based solely on shared valence cannot be explained by extant spreading activation models of semantic memory.

Separate affective and cognitive processing systems. That evaluative priming may involve the operation of a separate affective processing system that operates by different rules than does nonaffective cognitive processing is not a new idea (see Kuhl, 1986; LeDoux, 1989; Murphy & Zajonc, 1993; Zajonc, 1980a). In fact, Osgood et al. (1957) concluded from the dominant role that evaluation played in accounting for semantic similarity ratings that evaluation was a special semantic feature, possibly accessed first and even in the absence of activation of other semantic features.

In a test of this possibility, Bargh, Litt, Pratto, and Spielman (1989) presented trait adjectives below the subject's idiosyncratic threshold of conscious awareness. On a given trial, subjects were asked to respond either whether the stimulus was good or bad in meaning, or which one of two possibilities was a synonym of the target. As the presentation durations of the stimuli dropped farther below the consciousness threshold on successive blocks of trials (i.e., presentation durations became ever shorter), subjects retained the ability to respond correctly to the evaluative question at better than chance levels, but were not able to answer the semantic question with a greater than chance probability. Thus subjects had access to the evaluative information about the target independently of conscious access to the semantic meaning of the target. Path analyses also showed that the ability to answer the evaluation question was independent of the ability to answer the semantic question about that target, at least until performance on all questions dropped to chance levels at the briefest presentation durations.

Animal research by LeDoux and colleagues (e.g., 1986; LeDoux, Sakaguchi, & Reis, 1984) on the neural substrates of affective reactions to stimuli provides further support for the Osgood hypothesis. Brain areas (thalamic relay nuclei) representing the sensory features of acoustic stimuli associated with pain were found to develop subcortical synaptical connections directly to the amygdala (responsible for emotional responses to stimuli), bypassing the sensory cortex (LeDoux et al., 1984). In other words, associative connections developed directly from the stimulus feature representation to the affective response mechanism, bypassing the cognitive processing mechanism entirely. Given this combined evidence that affective pathways in the brain can be independent of cognitive analysis in animals, and that evaluative responses of human subjects can be independent of access to semantic meaning, it seems quite reasonable that automatic evaluation responses may be computed by an independent mental system that operates on somewhat different principles than nonevaluative automatic responses.

Finally, Murphy and Zajonc (1993) conducted several experiments to test their *affect primacy* hypothesis—the conjecture, highly consistent with the above discussion, that affective information is processed by an immediate mental system that is nonconscious. In support of this hypothesis, subliminally presented faces of positive versus negative emotional states were shown to influence judgments of the valence of ambiguous stimuli (Chinese ideographs), whereas supraliminal presentations of the same faces did not (see also Niedenthal, 1990). Moreover, in other experiments the valence of the subliminal prime was shown to affect judgments of valence, whereas other, nonaffective features of the prime (e.g., its size) were not shown to influence judgments about the target along those dimensions. Thus, affective information was processed immediately and influenced subsequent judgments whereas nonaffective information about the stimulus did not have this effect, and across several experiments, adding conditions in which the prime was processed consciously or strategically *reduced* or *eliminated* these affective priming effects.

The functional nature of initial evaluative screening. In addition to the question of mechanism, there is the question of why people automatically evaluate all stimuli that they come in contact with, no matter how mundane. There would seem to be some adaptive purpose served by screening all objects, people, and events in terms of their valence. Pratto and John (1991) have argued from their findings of greater attention allocated to negative trait stimuli for a mechanism of "automatic vigilance" toward negative stimuli, and Hansen and Hansen (1988) drew similar conclusions given their subjects' more efficient processing of angry versus happy faces relative to other faces in a set.

Also, from very different research areas, both Lazarus (1991) and LeDoux (1989) have concluded that stimuli are immediately and preconsciously evaluated or appraised in terms of their implications for the self. Indeed, LeDoux (1989) distinguishes between "cognitive" and "affective" processing on this basis: cognitive processing deals with the meaning of the stimulus *per se*, whereas affective processing deals with the implications of the stimulus for the self (see also Zajonc, 1980b). Our evidence of a ubiquitous automatic and preconscious evaluation effect is in line with models of emotion production that posit a primary stage of environmental appraisal (e.g., Lazarus, 1982, 1991, pp. 152–170; Ortony, Clore & Collins, 1988). The automatic classification of events as good or bad may be a first step to an emotional experience, as such appraisal models assume.

However, an important distinction is that according to these models, events are evaluated as to their self-relevance; that is, in terms of their relation to the person's important current goals. In this regard, we endorse the distinction made by Ortony et al. (1988) between different *targets* of evaluation or appraisal: whether an event, a person, or an object is being evaluated. According to this view, events are appraised in terms of one's goals, people are appraised according to personal standards of "praiseworthy" behavior, and objects are evaluated as "appealing" or not. Perhaps the automatic evaluation effect is most closely aligned with this latter variety of appraisal.

The repeated finding that differences in the chronic accessibility of the object-evaluation association does not matter to the "mere" automatic evaluation effect (as in the present research), whereas it does matter when more extensive conscious processing of the experimental stimuli is involved (as in the Fazio et al., 1986, experiments and Bargh et al., 1992, Experiment 1) may therefore be explicable in terms of the greater involvement of the cognitive processing system in the latter case. We believe that this distinction between what happens on the mere presence of the attitude object stimulus, when all conscious and strategic aspects of the paradigm are eliminated, and what occurs when conscious and strategic evaluative processing is involved, is critical to understanding when evidence for moderation of the automaticity effect by attitude strength variables will and will not be obtained. In six successive experiments (Bargh et al., 1992, Experiments 2 and 3; Chaiken & Bargh, 1993; and the present experiments), as aspects of the original paradigm that support conscious and strategic evaluative processing (as opposed to the mere effect of the attitude object itself) have been removed, evidence for moderation of the effect by attitude strength has essentially disappeared.

In this light, it is noteworthy that the Neely (1977) category-priming study on which the Fazio et al. (1986) paradigm was patterned varied the associative strength between the category name–category member pairs in the design, and did *not* find any moderation of the automatic category priming effect as a function of associative strength. Unlike Fazio et al., however, Neely (1977) did not instruct his subjects to consciously process the prime word in any way, merely to respond as to whether the target was a word or nonword. He (p. 252) concluded that "the present experiment demonstrates that a category-name prime will facilitate the subsequent processing of its good and poor exemplars to the same degree, even when they are targets in a task in which the subject is not logically required to access information on category membership."

Warren conducted two category-priming experiments in which associative strength was varied as a factor. In the first (Warren, 1974), subjects were required to hold the prime word in memory and the results revealed greater priming effects for the strong versus the weak category associates. In the second study (Warren, 1977), subjects were instructed to *ignore* the priming stimuli, and the results revealed equivalent priming effects regardless of strength of association between prime and target. A subsequent study by Lorch (1982) did find stronger priming effects with stronger associative strength, but in line with the above pattern he had instructed his subjects to read each category prime and decide whether the target was a category member or not—in other words, only when elements of strategic, conscious processing were added to the paradigm did moderation emerge. Stanovich and West (1983) also degraded the stimulus display (by positioning an asterisk between each letter of the target stimulus word), increasing the amount of conscious work the subject needed to do to recognize the word, and with this additional element of strategic processing did obtain larger priming effects as a function of associative strength.

Thus, in the category priming as well as the affect primacy and automatic attitude research literatures, moderation of priming effects by associative strength occurs only when there is a degree of conscious, strategic processing of the priming stimulus, and not in the absence of such instructions to the subject.

Automatic Evaluation Is the General Case, Moderation the Special Case

Fazio (in his Addendum to Chaiken & Bargh, 1993, p. 765) argues that the failure to obtain significant moderation of the automaticity effect by attitude strength does not mean the moderation effect does not exist, because this "would be tantamount to accepting the null hypothesis." Yet for a theory to be scientific or empirical, according to Popper (1935/1959), "...it must be possible for [it] to be refuted by experience" (p. 41). A theory that posits an automatic attitude activation effect that varies in probability with the strength of the object-evaluation association (Fazio, 1986, 1993) must (and does) predict that the observance of the automaticity effect will be an increasing function of associative strength, and must take consistent evidence of non-moderation as a refutation (see Popper, 1935/1959, p. 86). Many studies (the present three experiments; Bargh et al., 1992, Experiments 2 and 3; Chaiken & Bargh, 1993) have now obtained a statistically equivalent automaticity effect for the weakest as well as the strongest attitudes in the experiment, just as have several others in the category priming domain (see previous paragraph).

In light of all available evidence, it would appear that the basic, unmoderated automaticity effect (i.e., the Prime Valence × Target Valence interaction) is the more general case, and the moderated effect (i.e., the Attitude Strength × Prime Valence × Target Valence interaction) is a special case. Moderation has been found to occur, albeit under more circumscribed (and theory-irrelevant) conditions than the simpler two-way interaction. Although much remains to be learned about the mechanism(s) responsible for both the unmoderated and moderated automaticity effect, we believe that existing evidence indicates that automatic attitude activation is best understood as a basic and general *automatic evaluation effect* that is not dependent upon sufficient attitude strength. Regardless, our efforts to eliminate any potential artifactual causes of the automatic evaluation effect, and its robustness in the face of those challenges, allow us to conclude that the evaluation effect is genuine and truly unconditional, though its detection can be moderated or interfered with by the overlay of strategic conscious processes. Exactly why and how the general effect occurs are, to us, exciting directions for further research.

REFERENCES

Aizen, I. (1996). The directive influence of attitudes on behavior. In P. M. Gollwitzer & J. A. Bargh (Eds.), *The psychology of action: Linking cognition and motivation to behavior* (pp. 385–403). New York: Guilford.

Balota, D. A. & Lorch, R. F., Jr. (1986). Depth of automatic spreading activation. Mediated priming effects in pronunciation but not in lexical decision. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **12,** 336–345.

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.

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., Chaiken, S., Govender, R., & Pratto, F. (1992). The generality of the automatic attitude activation effect. *Journal of Personality and Social Psychology*, **62,** 893–912.

Bargh, J. A., Litt, J., Pratto, F., & Spielman, L. A. (1989). On the preconscious evaluation of social stimuli. In A. F. Bennett & K. M. McConkey (Eds.), *Cognition in individual and social contexts* (pp. 357–370). Amsterdam: Elsevier/North-Holland.

Bargh, J. A., Lombardi, W. J., & Higgins, E. T. (1988). Automaticity of person × situation effects on impression formation: 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.

Bellezza, F. S., Greenwald, A. G., & Banaji, M. R. (1986). Words high and low in pleasantness as rated by male and female college students. *Behavior Research Methods, Instruments & Computers*, **18,** 299–303.

Box, G. E. P., Hunter, W. G., & Hunter, J. S. (1978). *Statistics for experimenters*. New York: Wiley.

Chaiken, S. & Bargh, J. A. (1993). Occurrence versus moderation of the automatic attitude activation effect: Reply to Fazio. *Journal of Personality and Social Psychology*, **64,** 759–765.

Collins, A. M. & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, **82,** 407–428.

Eagly, A. H. & Chaiken, S. (1993). *The psychology of attitudes*. Fort Worth, TX: 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.

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). A practical guide to the use of response latency in social psychological research. In C. Hendrick & M. S. Clark (Eds.), *Review of personality and social psychology* (Vol. 11, pp. 74–97). Newbury Park, CA: Sage.

Fazio, R. H. (1993). Variability in the likelihood of automatic attitude activation: Data reanalysis and commentary on Bargh, Chaiken, Govender & Pratto (1992). *Journal of Personality and Social Psychology*, **64,** 753–758.

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

Fazio, R. H., 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.

Gollwitzer, P. M. (1990). Action phases and mind-sets. In E. T. Higgins & R. M. Sorrentino (Eds.), *Handbook of motivation and cognition* (Vol. 2, pp. 53–92). New York: Guilford.

Gollwitzer, P. M., Heckhausen, H., & Steller, B. (1990). Deliberative and implemental mind-sets: Cognitive tuning toward congruous thoughts and information. *Journal of Personality and Social Psychology*, **59,** 1119–1127.

Greenwald, A. G., Klinger, M. R., & Liu, T. J. (1989). Unconscious processing of dichoptically masked words. *Memory & Cognition*, **17,** 35–47.

Hansen, C. H. & Hansen, R. D. (1988). Finding the face in the crowd: An anger superiority effect. *Journal of Personality and Social Psychology*, **54,** 917–924.

Hedges, L. V. & Olkin, I. (1985). *Statistical methods for meta-analysis.* San Diego, CA: Academic Press.

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

Higgins, E. T., & Bargh J. A. (1987). Social perception and social cognition. *Annual Review of Psychology*, **38,** 369–425.

Higgins, E. T., Bargh, J. A., & Lombardi, W. (1985). Nature of priming effects on categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **11,** 59–69.

Higgins, E. T. & King, G. A. (1981). Accessibility of social constructs: Information-processing consequences of individual and contextual variability. In N. Cantor & J. F. Kihlstrom (Eds.), *Personality, cognition, and social interaction* (pp. 69–122). Hillsdale, NJ: Erlbaum.

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.

Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology*, **13,** 141–154.

Krosnick, J. A., Boninger, D. S., Chuang, Y. C., Berent, M. K., & Carnot, C. G. (1993). Attitude strength: One construct or many related constructs? *Journal of Personality and Social Psychology*, **65,** 1132–1151.

Kuhl, J. (1986). Motivation and information processing: A new look at decision making, dynamic change, and action control. In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition* (pp. 404–434). New York: Guilford.

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

LeDoux, J. E. (1989). Cognitive-emotional interactions in the brain. *Cognition and Emotion*, **3,** 267–289.

LeDoux, J. E., Sakaguchi, A., & Reis, D. J. (1984). Subcortical efferent projections of the medial geniculate nucleus mediate emotional responses conditioned by acoustic stimuli. *Journal of Neuroscience*, **4,** 683–698.

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

Lorch, R. F., Jr. (1982). Priming and search processes in semantic memory: A test of three models of spreading activation. *Journal of Verbal Learning and Verbal Behavior*, **21,** 468–492.

Murphy, S. T. & Zajonc, R. B. (1993). Affect, cognition, and awareness: Affective priming with optimal and suboptimal stimulus exposures. *Journal of Personality and Social Psychology*, **64,** 723–739.

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,** 226–254.

Niedenthal, P. M. (1990). Implicit perception of affective information. *Journal of Experimental Social Psychology*, **26,** 505–527.

Ortony, A., Clore, G. L., & Collins, A. (1988). *The cognitive structure of emotions*. New York: Cambridge Univ. Press.

Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1958). *The measurement of meaning*. Urbana: University of Illinois Press.

Pomerantz, E. M., Chaiken, S., & Tordesillas, R. (1995). Attitude strength and resistance processes. *Journal of Personality and Social Psychology,* **69**, 408–419.

Popper, K. (1935/1959). *The logic of scientific discovery*. New York: Harper & Row. (Originally published in 1935 as *Logik der Forschung*.)

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.

Pratto, F. & John, O. P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Personality and Social Psychology*, **61,** 380–391.

Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin*, **114,** 510–532.

Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, **104,** 192–233.

Smith, E. R. & Branscombe, N. (1987). Procedurally mediated social inferences: The case of category accessibility effects. *Journal of Experimental Social Psychology*, **23,** 361–382.

Smith, E. R. & Lerner, M. (1986). Development of automatism of social judgments. *Journal of Personality and Social Psychology*, **50,** 246–259.

Srull, T. K. & Wyer, R. S., Jr. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. *Journal of Personality and Social Psychology*, **37,** 1660–1672.

Stanovich, K. E. & West, R. F. (1983). On priming by a sentence context. *Journal of Experimental Psychology: General*, **112,** 1–36.

Tulving, E. & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, **5,** 381–391.

Warren, R. E. (1974). Association, directionality, and stimulus encoding. *Journal of Experimental Psychology*, **102,** 151–158.

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

Winer, B. J. (1971). *Statistical principles in experimental design* (2nd ed.). New York: McGraw–Hill.

Wyer, R. S., Jr. & Srull, T. K. (1989). *Memory and cognition in its social context*. Hillsdale, NJ: Erlbaum.

Zajonc, R. B. (1980a). Feeling and thinking: Preferences need no inferences. *American Psychologist*, **35,** 151–175.

Zajonc, R. B. (1980b). Cognition and social cognition: A historical perspective. In L. Festinger (Ed.), *Retrospections on social psychology* (pp. 180–204). New York: Oxford Univ. Press.