Journal of Personality and Social Psychology
2001, Vol. 81, No. 5, 815–827

0022-3514/01/$5.00 DOI: 10.1037//0022-3514.81.5.815

Spontaneous Prejudice in Context: Variability in Automatically Activated Attitudes

Bernd Wittenbrink
University of Chicago

Charles M. Judd and Bernadette Park
University of Colorado

The goal of the research reported in this article was to examine whether automatic group attitudes and stereotypes, commonly thought to be fixed responses to a social category cue, are sensitive to changes in the situational context. Two experiments demonstrated such variability of automatic responses due to changes in the stimulus context. In Study 1 White participants' implicit attitudes toward Blacks varied as a result of exposure to either a positive (a family barbecue) or a negative (a gang incident) stereotypic situation. Study 2 demonstrated similar context effects under clearly automatic processing conditions. Here, the use of different background pictures (church interior vs. street corner) for Black and White face primes affected participants' racial attitudes as measured by a sequential priming task. Implications for the concept of automaticity in social cognition are discussed.

Only a decade ago, the first empirical investigations emerged on the possibility that group attitudes and stereotypes may influence people's social perceptions and behaviors in an automatic fashion, outside of the individual's control (Devine, 1989; Gaertner & McLaughlin, 1983). Since then, however, a substantial number of studies have documented that such effects can occur (Banaji & Greenwald, 1995; Banaji & Hardin, 1996; Bargh, Chen, & Burrows, 1996; Blair & Banaji, 1996; Chen & Bargh, 1997; Devine, Monteith, Zuwerink, & Elliot, 1991; Dovidio, Evans, & Tyler, 1986; Dovidio, Kawakami, Johnson, Johnson, & Howard, 1997; Fazio, Jackson, Dunton, & Williams, 1995; Gilbert & Hixon, 1991; Greenwald et al., 1998; Hense, Penner, & Nelson, 1995; Kawakami, Dion, & Dovidio, 1998; Lepore & Brown, 1997; Locke, MacLeod, & Walker, 1994; Macrae, Bodenhausen, Milne, Thorn, & Castelli, 1997; Macrae, Milne, & Bodenhausen, 1994; Macrae, Stangor, & Milne, 1994; Moskowitz, Gollwitzer, Wasel, & Schaal, 1999; Perdue & Gurtman, 1990; Spencer, Fein, Wolfe, Fong, & Dunn, 1998; Wittenbrink, Judd, & Park, 1997). Distinguishing automatic activation from a controlled and intentional search for memory contents (Shiffrin & Schneider, 1977), this work leaves little doubt that stereotypes and group attitudes may indeed be activated spontaneously from memory, without the perceiver's intent, merely triggered by exposure to a relevant stimulus cue in the environment. Such automatic activation occurs quickly, within a few hundred milliseconds after stimulus exposure. It requires only very limited cognitive resources and is not controllable by the perceiver. In fact, the perceiver often remains unaware of the activation and its subsequent influences on judgment and behavior.

It is this latter quality of automatic stereotyping and prejudice that, in all likelihood, is responsible for much of the attention that the topic has received in the past few years. Perhaps researchers continue to be fascinated by the sources of people's behaviors that remain unknown to them, as has been true of psychological inquiry since its inception. However, it is more likely that the pragmatic implications of such "unconscious" stereotyping and prejudice have motivated this research. These pragmatic implications are indeed significant. For example, unconscious activation of negative cultural stereotypes has the potential to lead well-intentioned perceivers to walk away with a prejudiced impression of their interaction partners (Devine, 1989). Worse yet, this activation could lead to behaviors on the part of the perceiver that are likely to be reciprocated in a stereotype-confirming manner, thus resulting in an automatic self-fulfilling prophecy (Chen & Bargh, 1997). Moreover, to the extent that one can assess an individual's tendency to show automatic prejudice, researchers have at their disposal a genuine "bona fide pipeline" to people's group attitudes and beliefs, a measurement instrument that is not marred by social demand characteristics, as are standard self-report questionnaire measures (Fazio et al., 1995). After all, if respondents remain unaware of their prejudiced responses, they have little opportunity to tailor these responses to comply with perceived social standards.

The Obligatory Nature of Automatic Responses

Given the significant implications for both research and applied settings, it is hardly surprising that the issue of automaticity has become one of the central topics in the literature on group attitudes and stereotyping (see Blair, 2001). In this work, one particular characteristic that is often attributed to automatic activation of

Bernd Wittenbrink, Center for Decision Research, University of Chicago; Charles M. Judd and Bernadette Park, Department of Psychology, University of Colorado.

This research was supported by the William S. Fishman Faculty Research Fund at the University of Chicago, Graduate School of Business, and by National Institute of Mental Health Grant R01 MH45049. We thank Reid Hastie for his many helpful comments on an earlier version of this article.

Correspondence concerning this article should be addressed to Bernd Wittenbrink, Center for Decision Research, Graduate School of Business, University of Chicago, 1101 East 58th Street, Chicago, Illinois 60637. Electronic mail may be sent to bernd.wittenbrink@gsb.uchicago.edu.

attitudes and group-related beliefs is that they are more or less fixed and stable responses to a group-relevant stimulus cue. That is, given the passive nature of automatic activation, the fact that it is triggered by an external stimulus cue, and that it is not under the perceiver's volitional control, activation of certain memory contents is thought to be obligatory: It always and unconditionally follows exposure to the particular stimulus cue (see Bargh, 1999). Consequently, attitude activation has been described as a "reflexive" response following exposure to a relevant stimulus in the environment (Bargh, 1997, p. 3). Likewise, stereotypes have been thought to "be activated reflexively upon the mere presence" (Chen & Bargh, 1997, p. 546) of features that are diagnostic of a given social group (e.g., skin color, gender features; Bargh, 1994, 1999; Brewer, 1988). Thus, although one's overt responses may actually vary across different situations, seeing, for example, a Black person's face will always trigger memory activation of attributes that are stereotypically associated with Blacks (e.g., cheerful, poor).

There is, however, some evidence to suggest that spontaneous activation of group attitudes and stereotypes is not quite as reflexive as it has been assumed to be and that, in fact, such activation is not an obligatory response to a particular category cue. Most notably, Gilbert and Hixon (1991) showed that stereotype activation does not inevitably, in a reflexive manner, follow exposure to a stereotype target but that it depends on the availability of cognitive resources. Whereas Gilbert and Hixon's study used Asian targets, similar effects have also been reported by Spencer and his colleagues (Spencer et al., 1998) for the stereotyping of Blacks. Moreover, Blair and Banaji (1996) observed in a common sequential priming paradigm that spontaneous stereotype activation is dependent on participants' expectations about the relationship between prime and target stimuli. When participants expected the primes to be followed systematically by a counterstereotypic target, results no longer showed evidence of automatic activation of stereotypic associations. Finally, Macrae and his colleagues (Macrae et al., 1997) demonstrated that spontaneous stereotype activation in such priming experiments depends on the task under which participants encounter the priming stimuli. Processing the primes simply for their (conceptually irrelevant) surface features eliminated the occurrence of spontaneous stereotype activation.

Although these studies have not been without criticism (e.g., Bargh, 1999), to us the argument that exposure to a category cue does not always lead to (automatic) activation of the attitude or stereotype seems fairly plausible. In fact, our proposition in this article goes even further. We suggest that variations in the stimulus context affect not only whether a stimulus cue triggers activation of group-related memory contents but also what particular aspects of those contents it spontaneously activates. After all, what does it really mean to activate one's attitude toward, for example, Germans? What exactly is that attitude? Is it the negative attitude associated with German Nazis, the positive attitude associated with products made in Germany, or one's evaluation of those somewhat esoteric German psychologists? The notion that attitudes reflect a unified, solitary construct is, to say the least, highly disputed. Instead, attitudes are frequently conceptualized as multifaceted representations in memory (Schwarz & Strack, 1991; Tesser, 1978; Tourangeau & Rasinski, 1988; Wilson & Hodges, 1992; Wilson, Lindsey, & Schooler, 2000; Zanna & Rempel, 1988). Similarly, stereotypes also are likely to include a multitude of often contradictory attributes (e.g., Devine & Baker, 1991).

With regard to explicitly expressed attitudes, there is ample evidence that depending on what aspects of such a multifaceted representation become salient, different "attitudes" will emerge (Salancik & Conway, 1975). For example, questions regarding one's attitudes toward affirmative action are likely to yield different responses when they are framed by issues related to racial prejudice than when they are placed within the context of equal opportunity (cf. Kinder & Sanders, 1990). Likewise, to the extent that social contexts make salient different stereotypic attributes, one will observe different consequences of stereotype application (e.g., Bodenhausen, Schwarz, Bless, & Wänke, 1995).

We believe that it is rather likely that such context effects not only affect activation under controlled processing conditions but also affect activation that is not under the perceiver's voluntary control. For example, sitting among the crowd at the United Center in Chicago and being flashed with references to the group of Blacks, even the self-declared bigot might be more likely to spontaneously activate memory contents such as "Michael Jordan," "cheerful," and "athletic," rather than "Willie Horton," "lazy," or "poor." We suspect that the same bigot would show largely the opposite pattern of spontaneous activation were he or she ever to set foot in a primarily Black, poverty-stricken neighborhood on Chicago's Southside on a dark night. In both examples, category cues would yield activation of stereotypic memory contents, but, obviously, they would activate rather different aspects of the Black stereotype.

The argument that situational context may influence the outcome of cognitive processes that occur automatically also finds support in cognitive research. Here, evidence suggests that even basic perceptual processes are not as unconditionally linked to a specific stimulus input as was initially thought and that they are indeed quite malleable—albeit resource efficient and generally unconscious and uncontrollable (Kahneman, & Treisman, 1984).

For example, the processes by which we understand uttered sounds as speech meet all common criteria for automaticity. Readers will agree that, under most circumstances, they listen to their counterparts without constantly trying to figure out whether the person just uttered a *d* or a *t*. Identification of auditory input as a given speech pattern is indeed effortless and resource efficient and can occur involuntarily, without the perceiver's active control (Shiffrin, Pisoni, & Castaneda-Mendez, 1974). Yet, despite the fact that people carry out these identification processes without their control, there is nevertheless good evidence that these processes are not triggered in an unconditional fashion by a specified auditory stimulus. Instead, the execution of these identification processes is dependent on allocation of attentional resources (Nusbaum & Schwab, 1986) and the perceiver's expectations about the nature of the encountered stimulus: The same auditory stimulus may be heard as a portion of uttered speech or as birds chirping, depending on the perceivers' prior expectations (Remez, Rubin, Pisoni, & Carrell, 1981).

Similarly, looking at the pattern depicted in Figure 1, we immediately "see" the two words *THE CAT*. That is, although the shape depicting the middle letter in each word is identical, it may be seen as an *H* in one context and an *A* in another (Selfridge, 1955). Again, perception of the target stimulus is in both cases spontaneous, fast, and resource efficient. But the exact outcome of

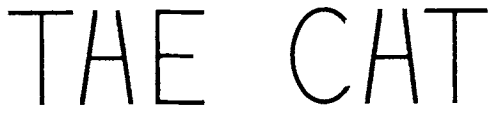

Figure 1. Context dependency in visual perception. The same shape may be "seen" as an *H* in one context and as an *A* in another. Reprinted from "Pattern Recognition and Modern Computers," by O. G. Selfridge, 1955, in *Proceedings of the 1955 Western Joint Computer Conference: Published by the Institute of Radio Engineers* (pp. 91–93), Los Angeles: Institute of Radio Engineers. Copyright 1955 by the Institute of Radio Engineers (now IEEE). Reprinted with permission.

this (by most definitions) automatic process depends on the context in which the target is encountered. In this latter example, variation in context does not consist of variation in the perceiver's expectations but of variation in the stimuli in which the target is embedded (i.e., *T* and *E* vs. *C* and *T*).

We believe that similar principles should also apply to the processing of social stimuli. Automatic attitudes and stereotypes should not be linked in an all-or-none fashion to a given category cue but should depend on the context in which the perceiver encounters that cue.

In this article, we report two experiments that illustrate such context effects on the memory contents activated spontaneously by a social category cue. These experiments varied the nature of additional stimuli in which the category cues were embedded. Activation of group attitudes and stereotypes then was assessed using two different procedures that have been commonly used in work on automatic attitudes and stereotyping, the Implicit Association Test (IAT; Greenwald et al., 1998) and the sequential priming paradigm (e.g., Dovidio & Gaertner, 1986; Fazio et al., 1995; Wittenbrink et al., 1997).

Study 1

Earlier, we speculated that category references to Blacks may trigger different memory contents depending on whether those references are encountered in a positive stereotypic context (e.g., a basketball arena) or a negative one (e.g., a poor urban neighborhood). Study 1 was intended to test exactly this conjecture. Half of the participants were shown a movie clip that depicted Black targets in a positive stereotypic situation, a family barbecue, whereas the remaining participants saw a movie clip with negative stereotypic context, a gang-related incident. We were interested in how this manipulation would affect the activation triggered by group references as measured by the IAT (Greenwald et al., 1998).

Method

Participants

Ninety-nine participants (18 Asian Americans, 10 African Americans, and 71 White Americans) were recruited on campus at the University of Chicago for paid participation ($10) in a 1.5 hr psychology experiment. They ranged in age from 17 to 37 years (*Mdn* = 20) and were predominantly undergraduate or graduate students. Participants who identified themselves as African American and 2 other participants who failed to follow instructions during the reaction time procedure were excluded from the data analyses, leaving a total of 87 participants (47 female, 40 male).

Procedures

The study was introduced to participants as an experiment on "how people tell stories" and, more specifically, on the role memory plays in the construction of story narratives. As part of these instructions, the experimenter explained that participants would watch a short movie excerpt and that it would be their task to write a story based on this excerpt. It was further explained that, later in the experiment and following a distraction task, participants would be asked questions about their stories and about the movie episode. The distraction task, in actuality the IAT, was introduced as a test of a person's ability to remain vigilant over a longer period of time. Participants were told that because individuals vary in how much cognitive energy they require to complete this test, the experiment would start out with a baseline assessment for the vigilance test. At the conclusion of these general instructions, the experimenter then mentioned in passing that in the event that they were to finish the experiment early, participants would be asked to complete a questionnaire for an unrelated study.

Following the introduction, participants were seated in front of a computer and told that the experimenter would now assess their baseline for the distraction task. Instructions for the IAT appeared on the computer screen. Participants read these instructions at their own pace and then proceeded with the baseline IAT. Half of the participants completed this baseline IAT with consistent response categories first, followed by the inconsistent response categories. For the remaining participants, this IAT order was reversed. During the experimental administration of this task, participants were seated at a distance of approximately 50 cm from the computer screen with their two index fingers positioned over the two response keys. Further details of the IAT trials are described below.

Once participants completed the first reaction time task, they were led to another room in the laboratory, where they were greeted by another experimenter and seated in front of a TV/VCR set. A shelf below the TV set and in clear sight of the participants held seven video tapes, labeled A through G. The experimenter explained that participants would now get to watch a short clip from one of the movies available. The experimenter then placed in the VCR one of the seven tapes, which for half of the participants contained a clip with a positive stereotype context and for the remaining participants contained a negative stereotype clip. After the short movie clip, participants had approximately 20 min to write an essay about the events depicted in the movie.

Following this story segment, participants returned to the computer room to participate in the alleged computer distraction task. It was explained that for them to better remember the movie episode while performing the computer task, they would from time to time see brief reminders of the movie clip. Participants then completed the experimental IAT. The actual clips that appeared at random intervals during these experimental IAT trials were matched with the movie excerpt participants had seen previously.

Immediately after they completed the experimental IAT, participants returned to the "video room" and filled out a questionnaire that included a series of questions relating to details of the movie episode (e.g., "List all protagonists that appeared in the movie. For each protagonist, give their name, describe their physical appearance etc."). Although the majority of these questions were included only for the sake of maintaining the experiment's cover story, the questionnaire included two items relevant to the actual purpose of the study. First, participants were asked to identify, among other features, the protagonists' race. We intended to use responses to this query as a manipulation check, assuring that participants correctly identified the movie targets' ethnicity. In addition, another question was included to determine whether participants were, as intended, unfamiliar with the movie from which the clip was taken. This question asked whether the participant knew the movie and, if so, to identify it.

Once participants had completed this questionnaire, the experimenter explained that there was still plenty of time left and that therefore they should please help out a friend of the experimenter's who was collecting data on an unrelated study. Participants then received a questionnaire containing six different explicit measures of racial attitudes. Specifically,

the questionnaire included an explicit measure commonly used to assess feelings toward social groups, the Feeling Thermometer rating scale, as well as a set of five belief-based attitude measures: the Modern Racism Scale by McConahay, Hardee, and Batts (1981), the Pro-Black and Anti-Black Scales of Katz and Hass (1988), and the Diversity and Discrimination Scales (both taken from Wittenbrink et al., 1997). To strengthen the reliability of the Thermometer scale, ratings for the two target groups (i.e., Blacks, Whites) were embedded in a series of filler target groups (e.g., Republicans, Democrats).

At the end of the experimental session, participants were debriefed about the actual purpose of the study, the nature of the IAT, and the potential influence that the movie clips were hypothesized to have on participants' IAT performance. They were then paid and dismissed.

Materials

Movie clips. Two movie excerpts were used to expose participants to either positive or negative stereotypic depictions of Blacks. Specifically, a 2-min segment from a feature movie entitled *Black & White & Red All Over* (Davis, McCoy, & Streeter, 1997) was used for the negative stereotype exposure condition. This segment depicted a gang-related incident in which a group of Black targets was seen arguing with each other, picking up a gun, and leaving the scene—apparently to confront an adversary. For the positive stereotype exposure condition, we used a 2-min segment of another feature movie, *Poetic Justice* (Singleton, 1993). The particular scene included in this segment showed a Black family in harmony together at an outdoor barbecue. Both movie segments included only relatively unknown actors and, on the basis of pretesting, were effective in eliciting narratives from participants that focused primarily on either positive or negative aspects of the stereotype.

IAT. Presentation of experimental stimuli and data collection was controlled by the PSYSCOPE software package (Version 1.2.4; Cohen, MacWhinney, Flatt, & Provost, 1993) on Apple Macintosh 7200/120 computers that were equipped with dedicated PSYSCOPE button boxes and 35-cm Apple color monitors. Stimuli were presented on a white background in black Geneva Macintosh font, 18 point, bold.

The procedure followed closely the original IAT described by Greenwald and his colleagues (Greenwald et al., 1998). In this task, participants are asked to categorize target words into one of two categories on the basis of a particular dimension of judgment (e.g., cancer—good vs. bad). The particular trials of interest are those that combine two types of judgments. For example, a series of trials may randomly present target words that consist of first names that are stereotypic of Blacks or Whites (e.g., Rasaan, Andrew) as well as nouns with either strong positive or strong negative evaluative connotation (e.g., love, cancer). The procedure's critical feature is that both category judgments (e.g., Black vs. White for names and good vs. bad for nouns) are made using only two response key (e.g., Black/good vs. White/bad). The time it takes to respond to target words in these critical trials, therefore, is thought to be influenced by the extent to which the two categories paired on a single key are associatively related in semantic memory. Faster responses should be observed for category combinations that are associatively related, whereas responses should slow down if the category combinations are inconsistent with the respondent's associations. Thus, in the case of combining social groups with evaluative categories, a spontaneous prejudice bias would be reflected in relatively faster responses on trials in which the out-group is paired with a negative category, relative to trials in which it is paired with a positive category. Conceptually, this prejudice bias is equivalent to what we have referred to as an indicator of "generalized prejudice" (Wittenbrink et al., 1997; Wittenbrink, Judd, & Park, 2001). It captures an individual's tendency to show a general negativity bias in associations with the out-group, independent of what the specific contents of these associations are.

Accordingly, the IAT that was of theoretical interest to us in the current study involved Black versus White judgments and good versus bad judgments. The target words consisted of 20 Black and 20 White first names and of 20 positive words and 20 negative words. The particular items used as targets were taken from Greenwald et al. (1998, Experiment 3).

In addition to the IAT trials of theoretical interest here, we also included a set of filler IAT trials to further disguise the actual purpose of the IAT procedure. For these filler IAT trials, participants had to categorize flowers and insects (e.g., daffodil—flower vs. insect), using 20 filler items that were also obtained from Greenwald et al. (1998, Experiment 1).

We organized the administration of IAT trials in blocks, varying judgments and key assignments between blocks. Each block presented participants in random order with the full set of target stimuli that were relevant to the block's judgment task. Specifically, the baseline IAT consisted of a sequence of five IAT blocks that make up the original Greenwald et al. (1998) procedure: (Blocks 1 and 2) separate practice trials for each individual judgment dimension, (Block 3) critical trials combining the two judgment dimensions, (Block 4) practice trials for the reverse evaluative judgments, and (Block 5) critical trials for the reverse combined judgments. These trials were followed by an additional block of filler IAT trials. The experimental IAT, administered after the movie manipulation, then repeated the critical IAT trial blocks for combined judgments, each of which were preceded by a block practicing the key assignment for the required evaluative judgments. In addition, because evidence from previous research suggests that the IAT effect can be affected by whether the consistent or the inconsistent trial blocks are presented first (Greenwald et al., 1998, Experiment 1), we counterbalanced this order across participants.

The complete baseline IAT consisted of a total of 360 trials (8 blocks per 40 trials). Of these, 80 trials were critical trials that were used for the assessment of the baseline IAT effect. The experimental IAT included another 240 trials, of which, again, 80 trials were critical for the subsequent data analyses.

As already mentioned, the experimental IAT also included displays of brief reminders of the movie clips that were used for the context manipulation. For this, we prepared five different 20-s excerpts from each of the original clips. The excerpts were then digitized so that they could be presented on the computer screen as part of the experimental IAT trials. Each of the five 20-s clips was displayed once during each of the two critical IAT blocks, interrupting the IAT procedure at a randomly chosen trial.

Results and Discussion

As expected, participants were unfamiliar with the movie clips used for the stereotype exposure manipulation. When asked about the movie title and actor names, all participants responded that they did not know the movie nor the actors. Also, all participants identified the movie protagonists correctly as African American.

IAT Response Latency Measure

Our primary objective in this study was to demonstrate that exposure to different aspects of a group stereotype affects spontaneous activation of group attitudes, as measured by the IAT. Thus, in reporting the results, we first focus on the analyses for the IAT response latency measure. For these analyses, we conducted a mixed model analysis of variance (ANOVA) that included two within-subject factors—assessment (baseline/experimental) and response assignment (consistent/inconsistent)—and two factors that varied between subjects—stereotype exposure (positive/negative) and order of IAT blocks (inconsistent first/consistent first).[1]

[1] Preliminary examination of the data revealed no significant effects involving participant gender. Consequently, this factor was not included in the final analyses.

Distributions for the response latency data show the common positive skew and a small number of outliers in which responses either were started prior to the actual target stimulus or were delayed because of temporary inattention (see Ratcliff, 1993). To address these problems in the present IAT data, we excluded response latencies faster than 300 ms and slower than 3,000 ms from the statistical analyses and then applied a log-transformation to the latencies.[2]

As explained earlier, in the IAT, relatively faster responses to consistent trials than inconsistent trials indicate a spontaneous prejudice bias, the "IAT effect." Figure 2 presents the average IAT effects observed at baseline and after the movie manipulation, separately for the two stereotype exposure conditions and retransformed into the millisecond metric.

Considering first the results for the baseline IAT, we find that the response latencies reveal an overall strong spontaneous prejudice bias (M = 136.70), $F(1, 86) = 136.80$, $p < .0001$. It seems, however, that participants in the positive and negative stereotype exposure conditions differed in the extent to which they displayed a spontaneous prejudice bias. Participants in the positive exposure condition showed a somewhat larger IAT effect than did those in the negative condition (Ms = 119.31 vs. 152.93). This difference, however, is not statistically significant, $F(1, 86) = 2.29$, $p = .134$, as it should not be because participants were assigned randomly to condition and the experimental procedure did not vary for the two conditions until after the time of the baseline assessment.[3]

The postmovie IAT assessment also yielded a sizable overall IAT effect (M = 62.60), $F(1, 86) = 68.44$, $p < .0001$. However, compared with the baseline assessment, this spontaneous prejudice bias was now reduced to less than half its original size. This overall reduction from Time 1 to Time 2 is statistically significant, $F(1, 86) = 25.34$, $p < .0001$, and most certainly reflects, in part, a practice effect on participants' ability to respond quickly to the target items and to do so in particular on the inconsistent trials.

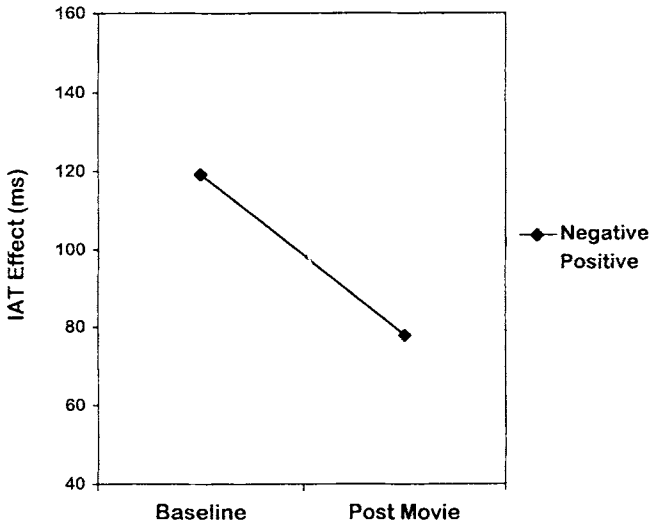

Figure 2. Implicit Association Test effect by stereotype condition and assessment (Study 2). Response time difference is in milliseconds between inconsistent and consistent judgment trials collapsed across both order conditions.

However, our main prediction for this study was that this change from Time 1 to Time 2 would be moderated by the stereotype exposure manipulation. In fact, comparison of the two exposure conditions shows that participants who were exposed to the positive group stereotype showed a significantly larger decrease in their IAT effect than did participants in the negative exposure condition (Ms = 104.53 vs. 41.50). This mean difference corresponds to the Assessment × Response Assignment × Stereotype Exposure three-way interaction, which proves to be significant, $F(1, 86) = 8.08$, $p = .006$. As suggested by the means, the three-way interaction effect emerges primarily because of the change in the IAT effect observed in the positive exposure condition: Whereas the change in the IAT effect from Time 1 to Time 2 is highly significant for participants in the positive exposure condition, $F(1, 86) = 30.79$, $p < .0001$, it is not significant for participants in the negative condition, $F(1, 86) = 2.43$, $p = .127$. In fact, participants in the positive exposure condition, who at baseline were the participant group with the larger spontaneous prejudice bias, actually showed a smaller IAT effect following the experimental manipulation than did participants in the negative exposure condition (Ms = 48.41 vs. 77.81). Even when we did not control for the baseline differences between the participant groups, this now-reversed postmovie difference in IAT effect was marginally significant, $F(1, 86) = 3.31$, $p = .072$.

Finally, the present latency data also show significant effects due to the order of the IAT blocks. Specifically, the overall IAT effect was reduced reliably when participants started with the inconsistent IAT trials rather than the consistent trials, $F(1, 86) = 10.06$, $p = .002$. A reverse order of IAT blocks also yielded an overall smaller change in the IAT effect from baseline to the postmovie assessment, $F(1, 86) = 4.84$, $p = .031$, and it did so more in the negative stereotype exposure condition than in the positive condition, $F(1, 86) = 5.55$, $p = .021$. Although these order effects are of limited theoretical interest here, the present data are further indication that counterbalancing the trial order seems essential for the IAT measure not to be confounded with this procedural variable.

Explicit Attitude Measures

In addition to the IAT measure, Study 1 also included the six explicit measures of racial attitudes. From responses to the belief-based attitude scales, we calculated for each scale participants' prejudice scores by recoding responses to reversed items and then averaging across all items of the scale. For the Feeling Thermometer measure, ratings for Blacks were subtracted from those for Whites. On all measures, higher scores indicate higher levels of out-group prejudice.

All five belief-based attitude scales had high internal consistency. As is common for a college student sample, they yielded what are, on average, relatively positive attitudes toward the target

[2] The same latency boundaries and data transformation were used by Greenwald et al. (1998), who, in addition, also recoded outlier responses using the boundary values. In the current data, the reported results are not affected in any substantive way by such data substitution.

[3] The experimenter responsible for this part of the experimental procedure (the IAT assessment) was unaware of participants' condition assignment.

Table 1
Intercorrelations Among Measures of Racial Prejudice (Study 1)

Measure	1	2	3	4	5	6	α	*M*
Explicit attitudes								
1. Modern Racism Scale	—						.82	1.71
2. Pro-Black Scale	.64***	—					.79	3.58
3. Anti-Black Scale	.58***	.32**	—				.79	2.73
4. Diversity Scale	.60***	.57***	.51***	—			.54	2.29
5. Discrimination Scale	.81***	.68***	.64***	.71***	—		.85	2.12
6. Feeling Thermometer	.43***	.40***	.29**	.36**	.34**	—		
Implicit Association Test								
Baseline effect	.18	.21*	.17	.21*	.21*	.31**		
Postmovie effect	−.03	.06	−.06	.09	.07	.21*		
Negative condition	.34*	.18	.08	.19	.36*	.17		
Positive condition	−.19	−.03	−.14	.01	−.03	.21		

Note. $N = 87$.
* $p \leq .05$. ** $p \leq .01$. *** $p \leq .0001$.

out-group. The correlations among these five scales show high coefficients and are in the expected direction. The Feeling Thermometer measure was also reliably correlated with the other explicit measures; however, as in other studies in which we assessed both belief-based attitudes and the Feeling Thermometer scale (Wittenbrink et al., 2001), it is so to a lesser extent. The top portion of Table 1 summarizes the internal consistency indices, scale means, and correlation coefficients for these measures.[4]

Our primary interest in collecting the explicit measures was to assess their relationship with participants' spontaneous group attitudes and, in particular, to determine how these relationships were affected by the context manipulation. To this end, we correlated the measures of participants' spontaneous prejudice bias, their IAT effects at baseline and postmovie, with their responses to the six attitude measures from the questionnaire (see bottom portion of Table 1).

Beginning with the correlations observed prior to the experimental manipulation, it is noticeable that the relationships between the explicit and automatic measures remain fairly weak ($rs \leq .21$, $n = 87$, $ps \geq .05$). The one exception is the Feeling Thermometer measure. On this measure, respondents state "how they feel" about the target group rather than indicate their agreement or disagreement with certain beliefs pertaining to the group. Likewise, the IAT is based on the strength with which evaluative labels (good, bad) rather than stereotypic attributes are associated with the target groups. The present correlation coefficients suggest that the Feeling Thermometer measure indeed taps into something similar to what is captured by this particular response time measure. Specifically, it shows a moderate but reliable relationship with the IAT effect ($r = .31$, $n = 87$, $p = .004$). Participants who indicated relatively more negative feelings toward the out-group on the Feeling Thermometer task also documented a stronger valence bias in the response time procedure. Their responses were also relatively faster in the condition in which out-group labels were paired with the label *bad* (see also Wittenbrink et al., 2001).

The second noteworthy point about the correlations between the explicit measures and the IAT effect is that they are clearly weakened by the experimental manipulation; this is primarily due to the effects of the positive stereotype exposure condition. That is, following the movie manipulation that exposed participants to the positive out-group stereotype, correlations between the IAT effect and the explicit measures were no longer systematic. Indeed, four correlation coefficients were in the direction opposite to what we expected. In contrast, for participants exposed to the negative stereotype, the relationship between the IAT measure and the explicit measure remained systematic, but the nature of this relationship was changed. In particular, two belief-based attitude measures, the Modern Racism Scale ($r = .34$, $n = 87$, $p = .03$) and the Discrimination Scale ($r = .36$, $n = 87$, $p = .018$), rather than the Feeling Thermometer measure ($r = .17$, $n = 87$, $p = .3$), showed moderately strong correlations with participants' postmovie spontaneous prejudice bias. Exposure to the negative out-group stereotype increased overlap on these measures and the IAT, whereas the relationship with the Feeling Thermometer measure was weakened. Participants who held more negative beliefs about the out-group were also more likely to show a spontaneous prejudice bias on the IAT.

The first experiment, therefore, provides good evidence that automatically activated group attitudes may vary with situational context. In this experiment, placement of the category cues into either a stereotypically positive or a stereotypically negative stimulus context reliably changed participants' responses on the IAT. When participants watched a video clip with a positive stereotypic context (i.e., a family barbecue), their responses showed a significantly larger decrease in spontaneous prejudice bias relative to the baseline assessment than when they saw a video with a negative stereotype context (i.e., a gang-related incident).

One important question raised by these findings concerns the role that intentional elaboration on the part of the perceiver plays in this variation of automatic attitudes. Study 1 included several reminders of the two alternative stereotype contexts during the attitude assessment (i.e., the experimental IAT blocks). However, prior to this assessment, participants first watched the video clips and then had the opportunity to elaborate on the different stereotype contexts. Thus, one interpretation of the current results is that

[4] All explicit attitude measures remained unaffected by the stereotype exposure manipulation, Pro-Black Attitudes, $F(1, 86) = 1.23$, $p = .271$; all others, $Fs \leq 0.34$, $ps \geq .36$.

the change in attitude activation was a result of participants' rumination about different aspects of their group-related beliefs and feelings and that such rumination temporarily increased the accessibility of different memory contents. This interpretation is consistent with other recent evidence suggesting that automatic activation of stereotypes may be subject to strategic interference—for example, when counterstereotypic expectations have been reinforced for a period of time (Blair & Banaji, 1996; Blair, Ma, & Lenton, 2001; Dasgupta & Greenwald, 2001).

Although we find it plausible that such intentional and strategic interference, practiced over a certain period of time, may produce variation in automatically activated attitudes, it seems less relevant for the kind of situations that we describe earlier in this article. For example, the pattern recognition effects described by Selfridge (1955) certainly do not require the perceiver to first elaborate on the various shapes of *H*s and *A*s that one may encounter. Similarly, we argue that the suspected differences in memory contents activated by a Black target at either the United Center in Chicago or the dark street corner do not result from prior elaboration about the different situations. Study 2, therefore, is intended to address this issue.

Study 2

Specifically, this second experiment involves a modified version of Fazio et al.'s (1995) evaluative priming paradigm. The paradigm is based on a sequential priming procedure in which participants are first primed with a group exemplar (e.g., a picture of a Black or a White target). The primes are followed by target adjectives, which vary in valence and which participants have to judge for their evaluative connotation (good/bad). Using this procedure, Fazio et al. (1995) showed that for White participants, Black faces facilitated responses to negative items more than to positive items. As was the case for the IAT, such a valence bias is again assumed to indicate activation of negative attitudes toward Blacks.

The current experiment essentially replicates this procedure, with one critical modification. The priming stimuli varied not only with respect to their group membership (Black/White) but also with regard to their context. Independent of group membership, half of the primes consisted of photos depicting the target person in a positive context stereotypic of Blacks (a church), whereas the remaining primes showed the target person in a negative stereotype context (a dilapidated street corner). Thus, different from Study 1, in this second experiment variation of situational context was an integral part of the response time procedure.

In addition to this primary change concerning the nature of the context manipulation, Study 2 addresses several other questions raised by the initial experiment. First, the design of Study 2 allows for within-subject comparisons between the positive and negative context conditions that could not be carried out for the data obtained in Study 1. Second, the use of entirely different stimulus materials in Study 2 enables us to verify that the context effects observed in the first experiment were not merely due to peculiarities of the two different movie clips used in the first experiment. In addition, by relying on a different response time procedure, we intended to illustrate the importance of context effects across multiple implicit measures of stereotyping and prejudice.[5]

Method

Overview

In this second experiment, participants first completed a computer-based reaction time procedure modeled after the Fazio et al. (1995) evaluative priming paradigm. Following this computer task, participants worked on an alleged filler task during which they completed the questionnaire with the six explicit racial attitude measures from Study 1.

Participants

Fifty participants (6 Asian Americans, 3 African Americans, 41 White Americans) were recruited in the same fashion as in Experiment 1 (median age = 19 years). Participants who identified themselves as African American were excluded from the data analyses, as was 1 other participant who failed to follow instructions during the reaction time procedure. Two additional participants were excluded from the analyses because they expressed suspicion about the purpose of the experiment. This resulted in a total of 44 participants (26 female, 18 male) who were included in the data analyses.

Materials

The reaction time task was presented using the same computer equipment that was used for Study 1. The priming stimuli for this task were based on 35 digital color photos of White and Black young adult males. All photos were head shots taken against a monochrome background. Using digital photo editing software, we removed the background from these photos and replaced it with the experimental stimulus contexts. These contexts consisted of an interior shot of a small Baptist church and the shot of a street corner with a large graffiti-covered wall. In other words, the original 35 photos were transformed into two sets of images that depicted the identical targets either standing in a church or standing at a street corner.

A set of 24 adjectives served as target items for the reaction time task. Whereas Fazio et al.'s (1995) stimulus materials included adjectives based solely on their evaluative connotation, the present study included items that varied both in valence and in their stereotypicality with regard to the two target groups (see Wittenbrink et al., 1997). Of these 24 adjectives, 8 were attributes that are stereotypic of Whites (e.g., educated, greedy), 8 were attributes that are stereotypic of Blacks (e.g., athletic, poor), and a final set of 8 items were nonstereotypic adjectives that are not part of the cultural stereotype for either of these two groups (e.g., appealing). In addition, item stereotypicality was crossed with item valence so that each set was made up of four positive and four negative adjectives. The Appendix presents all 24 target adjectives. During the experiment, these critical target items were complemented by an additional set of 16 filler items that were used on filler trials only. Presentation of all lexical stimuli occurred in black 18-point Times Macintosh font on a white background.

Procedures

Computer task. The reaction time task was introduced to participants as a study "on how people visually recognize words." It followed, in

[5] Study 2 also included an additional measure of out-group bias, a social distance measure assessing participants' choice of seating distance from an out-group target (e.g., Macrae, Bodenhausen, Milne, & Jetten, 1994). However, this measure yielded results that were difficult to interpret. Because the measure is only tangential to the primary goal of this research and because it was collected after all other parts of the procedure had been completed, we chose not to include it in this article. A description of procedure and results, however, is available on request from us.

principle, the evaluative priming procedure by Fazio et al. (1995) and was organized into four blocks of individually randomized response trials.

The first block assessed baseline response times for participants' evaluative judgments (good/bad) of the target items, presented without a group prime. Each trial in this block began with a cross mark (+) in the center of the computer screen. The cross mark was then replaced by the target item, which remained on the screen for 250 ms. Two seconds after a response had been made, the block continued with the next trial, until participants had seen all 24 target items plus 16 filler items and a set of 10 initial practice trials. For each target, we recorded response and latency of the response, measured from the target onset. These response latencies served as the participant's baseline latency in the later assessment of priming effects.

In keeping with the cover story used by Fazio et al. (1995), the second trial block consisted of an alleged "face-learning" task, in which participants had to attend to facial photos for the purpose of later identifying them from among several foils. Thus, in this second block participants saw a series of eight practice photos, each presented twice, once with the church background and once with the street background. Six of these photos depicted White targets, and two depicted Black targets. None of the photos was used again in the critical (fourth) priming block.

The third block involved a recognition test of the faces participants had just seen. Participants were shown a series of 16 photos (the 8 practice targets, plus 8 foils). Each target was shown with a neutral background and remained on the screen until the participant had pressed one of two keys, labeled *yes* and *no*.

The fourth trial block involved the critical sequential priming procedure. Participants were told that the two previous tasks, the word task and the face-learning task, would be combined to determine how much the distraction from the face-learning task would slow down performance on the word task. In addition, participants were led to believe that another face recognition test for the faces included in this block would follow later in the experiment. In actuality, this second recognition test was never administered, as it served only as a cover story.

The procedure for this fourth experimental block was the same as in the first block, with the exception that the cross mark was replaced by a priming sequence. For this priming sequence, each trial began with the display of a forward mask. The nature of this mask varied with context condition and consisted of the background pictures that were also used for the actual priming stimuli (i.e., the church interior and the street corner). After 1 s, this mask was replaced with the prime. Because the mask and the prime background were identical, the prime display left the impression that a person suddenly appeared in the scene. Presentation of the prime lasted 128 ms, followed by another 128-ms interval during which the screen remained blank, and then was followed by the target item.[6]

After an initial set of practice trials with filler photos and items, participants saw a total of 144 trials. Of these, 96 trials were experimental trials that fully crossed, within participants, the four factors of the study's design: prime (Black/White), prime context (positive/negative stereotype context), target item (stereotypic of Blacks/stereotypic of Whites/nonstereotypic), and valence (positive/negative). In other words, every participant saw all 24 target items four times, each paired once with one of the four different primes (Black/positive context, Black/negative context, White/positive context, White/negative context). Four different Black and four different White faces were used as primes. They were combined with the target items so that each face was used twice as a prime for each item type (i.e., Black stereotypic/positive valence)—once in the positive context and once in the negative context. Although the pairing of target items and faces was randomized across all participants, the same face was paired with the same target item in both context conditions.

In addition to these critical trials, participants saw another 48 filler trials on which an additional eight White faces were paired with filler items.

Additional measures. Once participants had completed the response time task, they were instructed that the final face recognition test would take place after a short break and were asked whether they would mind filling up the time by participating in another short experiment that was currently being conducted by a friend of the experimenter. All participants readily complied with this request. The questionnaire that participants filled out as part of this "experiment" contained the six explicit prejudice measures already used in Study 1.

At the end of the experimental session, the experimenter probed participants for suspicion about any aspects of the procedure and then debriefed participants about the actual purpose of the experiment. Finally, participants were thanked and paid for their participation.

Results and Discussion

The primary objective of Study 2 was to demonstrate that context-dependent variation in the spontaneous activation of group attitudes, as found in Study 1, can be observed even without prior elaboration about the nature of this context. In reporting the results, we again begin with the analyses for the response latency measure.

Response Latency Measure

Again, the latency data included a small number of outliers that were due to unintended responses or temporary inattention. As in Study 1, we excluded extremely fast and slow responses from the data analysis (responses faster than 150 ms and slower than 2 standard deviations above the individual participant's mean response time; 3.13% of all responses). The response latencies were then subjected to a log-transformation to approximate the data to a normal distribution. Next, participants' log-transformed response latencies from the experimental block were subtracted from their log-transformed baseline responses from Block 1 to determine the effect of the group primes on the response latencies. Thus, more positive values for these difference scores indicate greater response facilitation due to a group prime. They were analyzed as a function of four within-subject factors: (a) prime (Black/White), (b) prime context (positive/negative), (c) item stereotypicality (Black stereotypic/White stereotypic/nonstereotypic), and (d) item valence (positive/negative), collapsing across the four individual target items within each of the Stereotypicality × Valence cells of the design.[7] Mean facilitation scores, for ease of interpretation again retransformed into milliseconds, are given in Table 2.

Our central prediction for these response latency data was that the context manipulation would alter what responses were facilitated by the different group primes. A comparison of the mean facilitation scores in the top and bottom panels of Table 2 readily confirms this prediction. In the negative context condition, when the group primes were presented as part of an urban street scene, the facilitation scores of our participants reflect the kind of spontaneous prejudice bias reported in the original experiment by Fazio et al. (1995): Black faces disproportionately facilitated responses to negative target items. This effect is equivalent to the baseline

[6] Thus, the interval between prime onset and target onset (stimulus onset asynchrony, or SOA) was a total of 256 ms, significantly shorter than the SOA used by Fazio et al. (1995). We chose this more rigorous stimulus timing to eliminate any potential influences stemming from controlled processing of the priming stimuli (see Neely, 1977).

[7] As in Study 1, preliminary analyses of these scores revealed no significant effects of participant gender.

Table 2
Mean Response Facilitation by Prime Context in Milliseconds (Study 2)

	Prime type					
	Black faces			White faces		
	Item stereotypicality			Item stereotypicality		
Item valence	AA	NON	WA	AA	NON	WA
	Negative Prime Context					
Positive	−34.16	−11.18	−10.41	−1.39	−0.56	−1.30
Negative	66.81	56.29	21.14	5.55	4.70	3.16
	Positive Prime Context					
Positive	61.54	26.53	44.65	41.58	26.70	42.35
Negative	10.16	3.52	−6.33	−7.48	−6.51	−3.57

Note. AA = target items stereotypic of African Americans; NON = target items stereotypic of neither group; WA = target items stereotypic of White Americans.

IAT effect from Study 1 and has been documented by various researchers (e.g., Banaji & Hardin, 1996; Chen & Bargh, 1997; Dovidio et al., 1997; Fazio et al., 1995, Greenwald, et al. 1998; Kawakami et al., 1998; Lepore & Brown, 1997; Moskowitz et al., 1999; Wittenbrink et al., 1997). In contrast, the pattern of means observed in the positive context condition differed substantially from this by now familiar finding. In this condition, the mean facilitation scores showed absolutely no indication of a prejudiced valence bias. Instead, the same facial primes, now displayed as part of the church scene, yielded generally stronger facilitation for positive rather than negative target items.

Analyses of the response latency data confirm this predicted Context × Prime × Item Valence interaction, $F(1, 43) = 10.26$, $p = .003$. The analyses further reveal two significant lower order effects related to this three-way interaction: Context × Valence, $F(1, 43) = 36.30$, $p < .0001$, and context, $F(1, 43) = 4.06$, $p = .050$; all other Fs ≤ 1.62. Specifically, the two-way interaction refers to the fact that, overall, there was a lowered facilitation advantage for negative items in the positive context condition (Ms = −42.26 vs. 36.11), whereas the context main effect indicates that the positive context condition yielded stronger facilitation overall than did the negative condition (Ms = 19.43 vs. 8.22).[8]

Considering the results for the two context conditions separately, we find that facilitation scores for the negative context condition show an overall valence effect with relatively higher facilitation for negative than for positive target items (Ms = 26.27 vs. −9.83), $F(1, 43) = 7.71$, $p = .008$. However, this overall valence effect is attributable almost entirely to the already mentioned differential effect from Black face primes, which facilitated responses more when the target items were negative than when they were positive (Ms = 48.08 vs. −18.58). In contrast, facilitation from White faces was essentially the same for negative and positive items (Ms = 4.47 vs. −1.09). Thus, the overall valence effect was moderated by a significant Prime × Valence interaction, $F(1, 43) = 12.28$, $p = .001$. Additional tests involving just the Black face primes indicate that this facilitatory bias for Black faces for responses to negative items is observed across all three kinds of adjectives: (Black) stereotypic, nonstereotypic, and (White) counterstereotypic, $F(1,43) = 25.74$, $p < .0001$. However, the bias was somewhat stronger for adjectives that were stereotypic of the group prime than for the other items, $F(1, 43) = 5.16$, $p = .028$. No other significant effects were revealed by the analyses for this condition (all Fs ≤ 1.18).

As we have already explained, in the positive context condition, the effect of target item valence was reversed. The same facial primes yielded, overall, stronger facilitation for positive target items. Moreover, this was true for both Black face primes and White face primes. Indeed, the only significant effect that emerged from the analyses for this condition was a main effect due to item valence, $F(1, 43) = 11.91$, $p = .001$ (all other Fs ≤ 1.35).

Additional Measures

In addition to the response latency measure, the experiment also included the six questionnaire-based attitude scales from Study 1. The prejudice scores for each of the six questionnaire measures were calculated in the same fashion as in the previous study. The overall results are essentially the same as in Study 1, with one exception: The Anti-Black attitude measure showed uncharacteristically low internal consistency ($\alpha = .36$).

As in Study 1, our primary reason for collecting the explicit measures was to determine how their relationship with participants' automatic attitudes was affected by the context manipulation. Thus, separately for each context condition, we correlated participants' explicit attitude scores with their spontaneous prejudice indices from the response time data. As in Study 1, the resulting correlations remained weak (rs $\leq .17$, $n = 44$, ps $\geq .26$). And, similar to the results from Study 1, this is again especially true for the correlations involving the response bias assessed in a positive context (rs $\leq .12$, $n = 44$, ps $\geq .41$). Once again, the Feeling Thermometer measure was the exception to this pattern. At least for the negative context condition, the spontaneous prejudice index showed a correlation of $r = .35$, $n = 44$, $p = .02$ (positive context: $r = .05$, $n = 44$, $p = .752$). Thus, with regard to these correlations, the Study 2 results essentially repeat the general pattern of findings from the first experiment.

General Discussion

Our primary goal in this research was to determine whether automatic group attitudes and stereotypes, commonly thought to be fixed and invariant, are in fact sensitive to changes in the situational context. The findings from both experiments clearly demonstrate such potential for variation in automatic responses to social category cues.

[8] These effects of the context manipulation were not limited to only those target items that were, by nature of the Black stereotype, most closely related to the two stereotypic contexts included in the study. For example, results remained essentially the same even when items that are associated with the church context (religious, musical) were excluded from the analyses. Specifically, the critical Context × Prime × Item Valence interaction remained significant, although at a slightly lower level, given the reduced number of trials included in these analyses, $F(1, 43) = 6.12$, $p = .017$.

In Study 1, placement of the category cues into either a stereotypically positive or a stereotypically negative stimulus context reliably changed participants' responses on the IAT. In Study 2, results from a rather different assessment procedure that used entirely different stimulus materials again showed that the content of automatically activated memory contents may vary across different situations. Moreover, Study 2, unlike Study 1, manipulated context as a within-subject factor. That is, in Study 1, between-contexts comparisons were carried out across different individuals. In this initial study, we therefore had to take into account individual differences that already existed prior to the experimental manipulation to determine the effect of our context manipulation on participants' automatic attitudes. In Study 2, however, we observed effects of context as a result of variation within the same participants, thus strengthening our argument that under different circumstances, the same person may automatically activate different attitudes in response to an attitude object.

Moreover, the particular stimulus timing as well as the nature of the context manipulation used in Study 2 lead us to believe that the observed context effects reflect variation in automatic activation that occurs without prior elaboration or strategic interference on the side of the perceiver. That is, in Study 2, context was varied at the level of individual response trials, for which the presentation order was fully randomized for each individual participant. In other words, response trials were not presented in blocks repeating the same context, which would have allowed participants to practice and elaborate on the nature of the situation. In addition, the design of Study 2 crossed the context variable with group prime, so that a given context appeared, randomly, with both in-group and out-group faces. Furthermore, the stimulus timing for primes and target stimuli minimized the potential for controlled processing of the primes, whereas the observed effects were critically determined by the nature of the primes. That is, for the negative context condition, Black primes facilitated responses to negative items, whereas White primes yielded no evidence of such effects. All of these factors in combination make it unlikely that the current results reflect effects that stem from participants' rumination over certain aspects of their group-related beliefs and feelings or that require participants' strategic interference for suppressing otherwise automatic responses. Instead, they illustrate that a group reference may spontaneously activate quite different memory contents depending on the particular circumstances under which it is encountered.

This is not to say, however, that other context irrelevant memory contents also associated with a triggering stimulus cue are not being activated at all. In fact, a more detailed analysis of the perceptual and cognitive processes underlying concept activation would probably suggest a sequential mechanism whereby a larger set of memory contents is initially activated and is subsequently filtered for its contextual relevance. For example, research on text comprehension (e.g., Burgess & Simpson, 1988; Gernsbacher, Varner, & Faust, 1990) indicates that word stimuli yield broad activation of context-relevant and irrelevant concepts immediately following exposure to the stimulus (less than 100 ms). However, activation for context-irrelevant concepts drops rapidly thereafter (within a few hundred milliseconds). Similar findings have been reported for activation resulting from nonlinguistic stimuli as well (e.g., Gernsbacher & Faust, 1991). In other words, consistent with our present findings, this research also indicates that activation, although it is initially broad, is context specific and that the mechanism that produces this specific activation occurs quickly and does not require any active control on the side of the perceiver or reader.

In addition, the present experiments offer data that were intended to link the context differences in out-group prejudice observed at the automatic level to other, explicit measures of group attitudes. With regard to these explicit measures, we observed in both experiments a similar pattern of weak to moderate relationships. Of course, given their correlational nature and the relatively small sample sizes for relationships of the magnitude that one might expect between such diverse measures, these results remain far from definitive. However, we believe that the additional measures included in the experiments illustrate that the differences in spontaneous prejudice introduced by the context manipulations covary in meaningful and predictable ways with other measures of out-group prejudice and, therefore, are not merely "blips" in the distribution of millisecond response latencies.

Components of Group Attitudes

One particular finding that is of interest is that in both studies the response time procedures yielded somewhat stronger relationships with the Feeling Thermometer measure than with the belief-based attitude scales. The Feeling Thermometer measure asks respondents about their feelings toward the target group rather than about their beliefs pertaining to the group. Likewise, the two response time procedures that we used in the current experiments are based on evaluative judgments rather than on judgments that focus on conceptual aspects of the target items (e.g., lexical decision tasks). Moreover, in Study 2 the variation of item stereotypicality affected the observed priming effects only to a modest degree (i.e., only for Black primes in the negative context condition). This is consistent with results we recently obtained in related work in which we systematically varied the nature of the judgment task in a sequential priming procedure (Wittenbrink et al., 2001). In this research we found that evaluative judgments yielded overall Prime × Valence effects similar to those observed in the present studies, whereas a lexical decision task yielded the out-group valence biases that interacted with the stereotypicality of the target item. Individual-differences measures based on these response time procedures also showed different relationships with explicit attitude measures. The evaluative priming task showed relatively stronger relationships with a Feeling Thermometer measure, whereas the lexical decision task resulted in stronger relationships with belief-based measures like the Modern Racism Scale. What this suggests, therefore, is that these alternative measures, belief-based measures on the one hand and evaluative measures on the other hand, whether at the explicit or implicit level, seem to tap into separate and partially independent components of group attitudes: conceptual, belief-based aspects of prejudice versus general feelings and affective responses toward the target group (see also Amodio, Harmon-Jones, & Devine, 2000).

Automaticity in Social Cognitive Functioning

We believe that the findings from the current two experiments have important implications for the nature of automatic processing

in social cognition. Several years ago, Bargh (1989, 1994) offered an important clarification to the then already extensive evidence for automaticity in social information processing. He pointed out that automaticity, rather than describing a discrete form of cognitive functioning, varies substantially with regard to the conditions that have to be met for the automatic process to take place. According to this analysis, automaticity comes in a variety of shapes and flavors, each defined by its particular set of conditions. For example, goal-dependent automaticity requires activation of a particular interaction goal prior to encountering a stimulus cue (e.g., self- vs. other impression formation goals; Bargh & Tota, 1988), whereas postconscious automaticity takes place following some form of prior processing that is conscious and aware (e.g., certain priming effects in impression formation; Banaji, Hardin, & Rothman, 1993). In addition, this analysis also suggests that automatic processing may vary with regard to the number of necessary conditions, thus defining a continuum ranging from processes that are conditioned solely on the presence of a relevant stimulus cue to processes that require a more complex combination of circumstances.

We find this analysis particularly useful, as it emphasizes the fact that all automatic processing is conditional and as it emphasizes the need for an analysis of what specific conditions determine a given automatic process. Naturally, such an analysis can only benefit from empirical investigations that vary the relevant variables systematically. For instance, the conditions that determine whether behaviors trigger automatic trait inferences can only be identified correctly to the extent that they have been considered as potential candidates in the first place. Indeed, a number of automatic effects that at first appeared to be conditioned by a single factor (i.e., the presence of a stimulus cue) have turned out to depend on a more complicated set of circumstances once researchers began looking for them (e.g., the perceiver's goal state in the case of trait inferences; Uleman & Moskowitz, 1994).

With regard to group attitudes and stereotypes, research has only recently begun to specify more precisely the conditions of their automatic activation. As we already mentioned, there is now growing evidence that the conditions for automatic stereotype and attitude activation are more complex than the mere exposure to a category cue. Availability of attentional resources (Gilbert & Hixon, 1991), the perceiver's processing goals (Macrae et al., 1997; Moskowitz et al., 1999), and strategic preferences (Blair & Banaji, 1996; Devine, Plant, Amodio, Harmon-Jones, & Vance, 2000) have been identified so far as conditional factors determining whether a category cue will result in automatic stereotyping or attitude activation.

The findings from the present two studies are consistent with this notion of conditional automaticity. In fact, the present data suggest that variation in aspects of the situation other than the category cue cannot only eliminate, hinder, or suppress the automatic process, as demonstrated by the previously mentioned research, but that it is a critical determinant for the exact nature of the automatic process itself. In the present studies the social category cues proved to be capable of automatically activating group-related memory contents under a variety of contexts. However, depending on the context and, thus, depending on the particular conditions under which a triggering stimulus cue was encountered, the particular memory contents that were activated varied significantly.

Although automatic responses to social category cues do not appear to be determined solely by the presence or absence of a triggering cue, the general framework of conditional and multifaceted automaticity still leaves this as a possibility for automatic responses to other kinds of stimuli. However—especially with regard to the kind of stimulus situations that are generally of interest to social psychologists—we are skeptical about the viability of this form of automaticity, which is conditional only on the perceiver's exposure to the stimulus cue. That is, we believe that mere features of a distal stimulus, like the "Blackness" of a target person, are unlikely to be the sufficient condition for a particular response. Mere features have to be translated into percepts, a process that is determined not just by the particular feature in question but also by a variety of contextual factors like the perceiver's focus of attention, other stimulus features present in the environment, and so forth. To the extent that there is variation in these other determinants, there is likely to be variation in the resulting (spontaneous) cognitive activation. The two experiments we report demonstrate such potential for variation in automatic responses with regard to a particular set of stimuli, references to racial groups. We suspect that the present findings are not limited to automatic group attitudes and stereotypes but apply to automatic responses more generally.

Of course, this is essentially the argument that, a long time ago, Gestalt psychologists put forth against the then dominant doctrine that psychological functioning could be explained by linking psychological responses to a specific external stimulus. Despite the fact that Gestalt psychologists emphasized the importance of consciousness and subjective experience, many of the phenomena they investigated required little active control on the side of the perceiver. For example, Wertheimer's (1923) work on perceptual organization focused on the experiential nature of perception, yet it nevertheless was concerned with cognitive processes that clearly qualify as automatic. As this work on perceptual organization shows, stimuli may be perceived as part of a grouped entity or as individual units, depending on the particular context in which they are placed. Because such construal processes must take place for any stimulus to acquire subjective reality, Gestalt psychologists argued that psychological functioning could not be explained on the basis of an analysis of the "objective" stimulus features alone (Köhler, 1947).

Ultimately, the arguments advanced by Wertheimer, Köhler, and their colleagues proved to be critical for a more complete understanding of human functioning as well as instrumental in defining the field of social psychology. The distinction between active controlled processes in social–cognitive functioning and those that are passive and that occur involuntarily without conscious experience has already proven to be enormously important. However, this advancement in social psychological theorizing should not lead us to commit once again what Köhler referred to as the "experience error": assuming that a percept is nothing more than a copy of the distal stimulus object (Köhler, 1947, p. 162). We should not ignore one of the primary lessons learned in earlier days: that distal stimulus features are an insufficient determinant of human experience and behavior.

References

Amodio, D. M., Harmon-Jones, E., & Devine, P. G. (2000). The emotional and physiological components of race bias: Individual differences in attention- and emotion-modulated startle eyeblink responses. Manuscript submitted for publication.

Banaji, M. R., & Greenwald, A. G. (1995). Implicit gender stereotyping in judgments of fame. *Journal of Personality and Social Psychology, 68,* 181–198.

Banaji, M. R., & Hardin, C. D. (1996). Automatic stereotyping. *Psychological Science, 7,* 136–141.

Banaji, M. R., Hardin, C., & Rothman, A. J. (1993). Implicit stereotyping in person judgment. *Journal of Personality and Social Psychology, 65,* 272–281.

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. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition* (Vol. 1, pp. 1–40). Hillsdale, NJ: Erlbaum Press.

Bargh, J. A. (1997). The automaticity of everyday life. In R. S. Wyer (Ed.), *The automaticity of everyday life: Advances in social cognition* (Vol. 10, pp. 1–61). Mahwah, NJ: Erlbaum Press.

Bargh, J. A. (1999). The cognitive monster: The case against the controllability of automatic stereotype effects. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 361–382). New York: Guilford Press.

Bargh, J. A., Chen, M., & Burrows, L. (1996). Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action. *Journal of Personality and Social Psychology, 71,* 230–244.

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.

Blair, I. V. (2001). Implicit stereotypes and prejudice. In G. B. Moskowitz (Ed.), *Cognitive social psychology: On the tenure and future of social cognition* (pp. 359–374). Mahwah, NJ: Erlbaum.

Blair, I. V., & Banaji, M. R. (1996). Automatic and controlled processes in stereotype priming. *Journal of Personality and Social Psychology, 70,* 1142–1163.

Blair, I. V., Ma, J. E., & Lenton, A. P. (2001). Imagining stereotypes away: The moderation of implicit stereotypes through mental imagery. *Journal of Personality and Social Psychology, 81,* 828–841.

Bodenhausen, G. V., Schwarz, N., Bless, H., & Wänke, M. (1995). Effects of atypical exemplars on racial beliefs: Enlightened racism or generalized appraisals? *Journal of Experimental Social Psychology, 31,* 48–63.

Brewer, M. B. (1988). A dual process model of impression formation. In T. K. Srull & R. S. Wyer (Eds.), *Advances in social cognition* (Vol. 1, pp. 1–36). Hillsdale, NJ: Erlbaum Press.

Burgess, C., & Simpson, G. B. (1988). Cerebral hemispheric mechanisms in the retrieval of ambiguous word meanings. *Brain & Language, 33,* 86–103.

Chen, M., & Bargh, J. A. (1997). Nonconscious behavioral confirmation processes: The self-fulfilling consequences of automatic stereotype activation. *Journal of Experimental Social Psychology, 33,* 541–560.

Cohen, J., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: An interactive graphical system for designing and controlling experiments in the psychology laboratory using Macintosh computers. *Behavior Research Methods, Instruments, and Computers, 25,* 257–271.

Dasgupta, N., & Greenwald, A. G. (2001). On the malleability of automatic attitudes: Combating automatic prejudice with images of admired and disliked individuals. *Journal of Personality and Social Psychology, 81,* 800–814.

Davis, D., McCoy, H., & Streeter, K. (Directors). (1997). *Black & white & red all over* [Motion picture]. United States: Spectrum Films.

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56,* 5–18.

Devine, P. G., & Baker, S. M. (1991). Measurement of racial stereotype subtyping. *Personality and Social Psychology Bulletin, 17,* 44–50.

Devine, P. G., Monteith, M. J., Zuwerink, J. R., & Elliot, A. J. (1991). Prejudice with and without compunction. *Journal of Personality and Social Psychology, 60,* 817–830.

Devine, P. G., Plant, E. A., Amodio, D. M., Harmon-Jones, E., & Vance, S. L. (2000). *Exploring the relationship between implicit and explicit prejudice: The role of motivations to respond without prejudice.* Unpublished manuscript, University of Wisconsin—Madison.

Dovidio, J. F., Evans, N., & Tyler, R. (1986). Racial stereotypes: The contents of their cognitive representations. *Journal of Experimental Social Psychology, 22,* 22–37.

Dovidio, J. F., & Gaertner, S. L. (1986). Prejudice, discrimination, and racism: Historical trends and contemporary approaches. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 1–34). Orlando, FL: Academic Press.

Dovidio, J. F., Kawakami, K., Johnson, C., Johnson, B., & Howard, A. (1997). On the nature of prejudice: Automatic and controlled processes. *Journal of Experimental Social Psychology, 33,* 510–540.

Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology, 69,* 1013–1027.

Gaertner, S. L., & McLaughlin, J. P. (1983). Racial stereotypes: Associations and ascriptions of positive and negative characteristics. *Social Psychology Quarterly, 46,* 23–30.

Gernsbacher, M. A., & Faust, M. E. (1991). The mechanism of suppression: A component of general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 245–262.

Gernsbacher, M. A., Varner, K. R., & Faust, M. E. (1990). Investigating differences in general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 430–445.

Gilbert, D. T., & Hixon, J. G. (1991). The trouble of thinking: Activation and application of stereotypic beliefs. *Journal of Personality and Social Psychology, 60,* 509–517.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology, 74,* 1464–1480.

Hense, R. L., Penner, L. A., & Nelson, D. L. (1995). Implicit memory for age stereotypes. *Social Cognition, 13,* 399–415.

Kahneman, D., & Treisman. A. (1984). Changing views of attention and automaticity. In R. Parasuraman & D. R. Davies (Eds.), *Varieties of attention* (pp. 29–61). San Diego, CA: Academic Press.

Katz, I., & Hass, R. G. (1988). Racial ambivalence and American value conflict: Correlational and priming studies of dual cognitive structures. *Journal of Personality and Social Psychology, 55,* 893–905.

Kawakami, K., Dion, K. L., & Dovidio, J. F. (1998). Racial prejudice and stereotype activation. *Personality and Social Psychology Bulletin, 24,* 407–416.

Kinder, D. R., & Sanders, L. M. (1990). Mimicking political debate with survey questions: The case of White opinion on affirmative action for Blacks. *Social Cognition, 8,* 73–103.

Köhler, W. (1947). *Gestalt psychology.* (Rev. ed.). New York: Liveright Publishing.

Lepore, L., & Brown, R. (1997). Category and stereotype activation: Is prejudice inevitable? *Journal of Personality and Social Psychology, 72,* 275–287.

Locke, V., MacLeod, C., & Walker, I. (1994). Automatic and controlled

activation of stereotypes: Individual differences associated with prejudice. *British Journal of Social Psychology, 33,* 29–46.

Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Jetten, J. (1994). Out of mind but back in sight: Stereotypes on the rebound. *Journal of Personality and Social Psychology, 67,* 808–817.

Macrae, C. N., Bodenhausen, G. V., Milne, A. B., Thorn, T. M. J., & Castelli, L. (1997). On the activation of social stereotypes: The moderating role of processing objectives. *Journal of Experimental Social Psychology, 33,* 471–489.

Macrae, C. N., Milne, A. B., & Bodenhausen, G. V. (1994). Stereotypes as energy-saving devices: A peek inside the cognitive toolbox. *Journal of Personality and Social Psychology, 66,* 37–47.

Macrae, C. N., Stangor, C., & Milne, A. B. (1994). Activating social stereotypes: A functional analysis. *Journal of Experimental Social Psychology, 30,* 370–389.

McConahay, J. B., Hardee, B. B., & Batts, V. (1981). Has racism declined in America? It depends on who is asking and what is asked. *Journal of Conflict Resolution, 25,* 563–579.

Moskowitz, G. B., Gollwitzer, P. M., Wasel, W., & Schaal, B. (1999). Preconscious control of stereotype activation through chronic egalitarian goals. *Journal of Personality and Social Psychology, 77,* 167–184.

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.

Nusbaum, H. C., & Schwab, E. C. (1986). The role of attention and active processing in speech perception. In E. C. Schwab & H. C. Nusbaum (Eds.), *Pattern recognition by humans and machines* (Vol. 1, pp. 113–157). San Diego, CA: Academic Press.

Perdue, C. W., & Gurtman, M. B. (1990). Evidence for the automaticity of ageism. *Journal of Experimental Social Psychology, 26,* 199–216.

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

Remez, R. E., Rubin, P. E., Pisoni, D. B., & Carrell, T. D. (1981). Speech perception without traditional speech cues. *Science, 212,* 947–950.

Salancik, G. R., & Conway, M. (1975). Attitude inferences from salient and relevant cognitive content about behavior. *Journal of Personality and Social Psychology, 32,* 829–840.

Schwarz, N., & Strack, F. (1991). Context effects in attitude surveys: Applying cognitive theory to social research. In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (Vol. 2, pp. 31–49). London: Wiley.

Selfridge, O. G. (1955). Pattern recognition and modern computers. *Proceedings of the 1955 Western Joint Computer Conference: Published by the Institute of Radio Engineers.* (pp. 91–93). Los Angeles: Institute of Radio Engineers.

Shiffrin, R. M., Pisoni, D. B., & Castaneda-Mendez, K. (1974). Is attention shared between the ears? *Cognitive Psychology, 6,* 190–215.

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

Singleton, J. (Writer/Director). (1993). *Poetic justice* [Motion picture]. United States: Columbia Pictures.

Spencer, S. J., Fein, S., Wolfe, C. T., Fong, C., & Dunn, M. A. (1998). Automatic activation of stereotypes: The role of self-image threat. *Personality and Social Psychology Bulletin, 24,* 1139–1152.

Tesser, A. (1978). Self-generated attitude change. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 11, pp. 289–338). New York: Academic Press.

Tourangeau, R., & Rasinski, K. A. (1988). Cognitive processes underlying context effects in attitude measurement. *Psychological Bulletin, 103,* 299–314.

Uleman, J. S., & Moskowitz, G. B. (1994). Unintended effects of goals on unintended inferences. *Journal of Personality and Social Psychology, 66,* 490–501.

Wertheimer, M. (1923). Untersuchungen zur Lehre der Gestalt, II [Laws of organization in perceptual forms]. *Psychologische Forschung, 4,* 301–350.

Wilson, T. D., & Hodges, S. D. (1992). Attitudes as temporary constructions. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments* (pp. 37–66). Hillsdale, NJ: Erlbaum.

Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review, 107,* 101–126.

Wittenbrink, B., Judd, C. M., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology, 72,* 262–274.

Wittenbrink, B., Judd, C. M., & Park, B. (2001). Evaluative versus conceptual judgments in automatic stereotyping and prejudice. *Journal of Experimental Social Psychology, 37,* 244–252.

Zanna, M. P., & Rempel, J. K. (1988). Attitudes: A new look at an old concept. In D. Bar-Tal & A. W. Kruglanski (Eds.), *The social psychology of knowledge* (pp. 315–334). Cambridge, England: Cambridge University Press.

Appendix

Target Items Used in Study 2

Items stereotypic of African Americans	Nonstereotypic items	Items stereotypic of White Americans
religious	desirable	intelligent
cheerful	fabulous	successful
athletic	pleasant	educated
musical	wonderful	responsible
poor	awful	boring
violent	horrible	uptight
lazy	repulsive	greedy
threatening	rotten	selfish

Received March 3, 2000
Revision received February 2, 2001
Accepted February 7, 2001 ■