Racial Prejudice and Stereotype Activation

Kerry Kawakami
Kenneth L. Dion
University of Toronto
John F. Dovidio
Colgate University

The main goal of the present study was to investigate the interrelationships among prejudice and the endorsement and activation of cultural stereotypes. The endorsement of stereotypes was examined using an attribution estimate measure. Stereotype activation was examined using a pronunciation task with short and long stimulus onset asynchronies (300 ms and 2000 ms) to induce automatic and controlled processing. As expected, high prejudiced participants endorsed cultural stereotypes to a greater extent than low prejudiced participants. Furthermore, for high prejudiced participants, Black category labels facilitated stereotype activation under automatic and controlled processing conditions. For low prejudiced participants, no evidence of differential activation was found for stereotypes relative to nonstereotypes as a function of category labels under either processing condition. In addition, stereotype activation was correlated with individual differences in stereotype endorsement.

Traditionally, stereotyping has been conceived as both a cause and a consequence of prejudice (Allport, 1954; Smith, 1993). Despite this theoretical link, research examining the assumed relationship between prejudice and stereotyping has often obtained weak or conflicting findings (Brigham, 1971; Dovidio, Brigham, Johnson, & Gaertner, 1996; Esses, Haddock, & Zanna, 1993; Mackie & Hamilton, 1993). In a meta-analytic review of the literature, Dovidio et al. (1996) found that although individual differences in stereotyping were statistically significantly related to prejudice overall, the effect size was modest ($r = +.25$) and highly variable across studies.

Recent research on stereotype activation suggests a possible moderator of the relationship between prejudice and stereotyping: the automaticity of the processing conditions. Devine (1989) proposed that the expression of stereotypes is conditional on (a) the characteristics of the processing environment (i.e., whether there is time for controlled processes to modify automatic effects) and (b) the prejudice level of the participants. With respect to racial stereotypes, for example, Devine (1989) hypothesized that because of common socialization experience, both low and high prejudiced Whites have the same group-based representations of Blacks[1] and automatically activate these representations when they encounter a member (or symbolic equivalent) of the group. When controlled processing is possible, however, low prejudiced Whites inhibit this tendency, whereas high prejudiced Whites continue to show this effect.

Supportive of this position, Devine (1989) found that when subliminally primed with Black stereotypic words, both high and low prejudiced participants interpreted ambiguous behavior in terms of the activated stereotypic schema (e.g., as hostile). Furthermore, as Devine (1989) proposed, although high and low prejudiced Whites showed similar patterns of stereotype activation under automatic processing conditions, they displayed different patterns under more controlled conditions. The self-reports of high prejudiced Whites revealed more evidence of negative personal stereotypes than did the self-reports of low prejudiced Whites. Devine (1989) concluded that whereas low prejudiced Whites consciously inhibit their initially biased responses, high prejudiced Whites do not.

Authors' Note: Jim Olson served as action editor for this manuscript. We are grateful for the helpful comments and suggestions offered by Jim Olson, Rupert Brown, Lorella Lepore, and the anonymous reviewers. This study was facilitated by a Social Sciences and Humanities Research Council (SSHRC) Award to the first author. Correspondence concerning this article should be addressed to Kerry Kawakami, Department of Social Psychology, University of Nijmegen, Postbus 9104, 6500 HE Nijmegen, the Netherlands; e-mail: kawakami@psych.kun.nl.

PSPB, Vol. 24 No. 4, April 1998 407-416

Conflicting evidence does exist, however, about whether people differing in prejudice have different degrees and types of automatic stereotype activation. Consistent with Devine's (1989) theorizing, Banaji and Greenwald (1995) found no systematic relationship between attitudes toward women and implicit sex stereotyping using a false-fame memory paradigm. Nevertheless, other data suggest that high and low prejudiced persons differentially display stereotype activation and may actually differ in their cognitive stereotypic representations for Blacks (Lepore & Brown, 1997; Wittenbrink, Judd, & Park, 1997) and for the elderly (Hense, Penner, & Nelson, 1995)—even in situations that do not permit controlled processing. People higher in prejudice displayed greater automatic stereotype activation of consensual stereotypes.

Given this conflicting evidence, one goal of the present study was to examine how individual differences in racial prejudice and personal stereotypes are related to the activation of Black stereotypes. Specifically, the present research attempted to reconcile the apparently conflicting findings by considering both methodological and theoretical issues. Methodologically, in the previous research the techniques used to assess controlled responses and automatic activation are quite different. For example, in Devine's (1989) work on the disassociation hypothesis, the controlled responses were self-reports of personal stereotypes, whereas the assessment of automatic activation was the degree to which impressions of another were influenced by the subliminal presentation of Black stereotypic words. In contrast, in the present research, activation of Black stereotypes was assessed in automatic and controlled processing conditions (varied as a function of stimulus onset asynchrony; Neely, 1977, 1991; see also Locke, MacLeod, & Walker, 1994) using the same type of response (pronunciation latencies) within a single paradigm. This strategy permits more direct comparisons that are uncomplicated by differences in response type.

Theoretically, the present work examined the role of personal endorsement of cultural stereotypes. Devine (1989) proposed that because of the repeated and virtually unavoidable exposure to pervasive cultural stereotypes, both high and low prejudiced individuals will automatically activate these representations when they are presented with representations of those groups regardless of their personal level of endorsement of these stereotypes (i.e., personal stereotypes). Recently, Lepore and Brown (1997) highlighted an important distinction between stereotype priming and category priming. Stereotype priming involves cueing stereotypic characteristics (e.g., lazy) directly, with or without the category label (e.g., Black). Categorization priming occurs when only the category (e.g., Black) is cued in isolation of any stereotypic characteristics. Lepore and Brown noted that Devine's (1989, Study 2) research involved both stereotypic priming and category priming simultaneously, and they observed, "Many primes had clear negative connotations . . . that could have directly cued hostility" (p. 276). The absence of differences in the responses of high and low prejudiced participants in the Devine study may thus have occurred because of the direct activation of semantic associations involved in stereotype priming rather than because of a close association between the category alone and the stereotype.

Lepore and Brown (1997) further argued that "high- and low-prejudice people's representations of the social group may not differ in terms of content (at least for stereotype knowledge) but stronger links may have developed for different characteristics" (p. 277). Lepore and Brown reasoned that, as a consequence of this differential strength of associative links with the category, high and low prejudiced people would show divergent automatic stereotype activation as a function of category priming. Consistent with their hypothesis, using Devine's (1989, Study 2) priming and subsequent impression formation procedure, Lepore and Brown found that when only the category was primed, high prejudiced participants showed evidence of automatic negative stereotype activation, whereas low prejudiced participants did not (and, in fact, tended to show activation of positive associates). When stereotype priming was involved, however, both high and low prejudiced participants demonstrated comparable levels of stereotype activation.

The present research further investigated this theoretical issue by directly examining the relation between personal *endorsement* of cultural stereotypes (shared beliefs about the characteristics possessed by members of a social group) and the activation of these stereotypes within a category priming paradigm. Because participants were given the time and opportunity to ascribe stereotypic traits deliberately to the particular categories, this process is considered to be controlled. Thus, the degree of the participants' endorsement of the cultural stereotypes was expected to vary as a function of prejudice (Devine, 1989; Esses et al., 1993; Lepore & Brown, 1997).

Participants' *activation* of cultural stereotypes, alternatively, was assessed with a word pronunciation task. Specifically, participants, who were classified as high or low in prejudice, were presented with a category prime (Black, White, or CCC [a neutral baseline]) followed by a positive or negative Black stereotypic target word or nonstereotypic target word. Their task was simply to pronounce the target word. Response latency was the dependent measure. A number of studies have revealed that this procedure may produce a particularly sensitive

measure of automatic processing because the paradigm does not foster task-specific strategies that can obscure the effects of automaticity (Balota & Chumbly, 1984; Balota & Lorch, 1986; Bargh, Chaiken, Raymond, & Hymes, 1996; Joordens & Besner, 1992; Ratcliff & McKoon, 1988). Conditions for automatic and controlled processing were induced in this paradigm by varying the length of time between the onset of the presentation of the prime and the presentation of the target word. Under short SOAs (stimulus onset asynchronies) (300 ms), responses were assumed to be automatic because the stimuli were presented too briefly for participants to engage, focus, and commit attention intentionally to their responses (Neely, 1977, 1991). Long SOAs (2,000 ms) were assumed to permit more controlled processes because participants had sufficient time to manipulate their responses strategically. In the present study, the repeated presentation of racial category primes and stereotypic traits makes the racial focus of the task salient, which normally produces motivations to respond in socially desirable, unbiased ways (Dovidio & Fazio, 1992), even in simple response latency tasks (Messick & Mackie, 1989).

If knowledge of cultural stereotypes is the primary determinant of automatic stereotype activation, as Devine (1989) posits, then high and low prejudiced participants would be expected to demonstrate equivalent levels of stereotype activation under short SOA conditions, but low prejudiced participants would suppress these responses under long SOA conditions. Within the present paradigm, support for Devine's position would be demonstrated by a four-way Prime × Stereotypicality × SOA × Prejudice interaction. In the 300 ms SOA condition, high and low prejudiced participants would be expected to show similar Prime × Stereotypicality interactions reflecting faster response times to Black stereotypic words (relative to nonstereotypic words) following the Black prime than following the White prime. In the 2,000 ms SOA condition, which permits controlled processes to operate, high prejudiced participants would be expected to display this pattern, whereas low prejudiced participants would inhibit their stereotypic responses and reveal no differential association as a function of the prime.

In contrast to Devine's position, Lepore and Brown (1997) propose that when category priming alone is involved, individual differences in prejudice, mediated through systematic differences in personal endorsement of stereotypes, will relate to automatic stereotype activation. This position suggests that because of their different mental representation of Blacks (a) high and low prejudiced people will differ in their endorsement of cultural stereotypes, and thus (b) high and low prejudiced people will display different patterns of automatic as well as controlled activation of stereotypes as a function of category primes. From this perspective, a three-way Prime × Stereotypicality × Prejudice interaction would be expected that was independent of the SOA condition. In particular, high prejudiced participants would be expected to show a Prime × Stereotypicality interaction reflecting stereotypic associations across both the 300 ms and 2,000 ms SOA conditions, whereas low prejudiced participants would be expected to show no evidence of differential stereotypic association within either SOA condition.

METHOD

Participants

Although 154 undergraduate students at the University of Toronto participated in the experiment to fulfill one option for their course requirements, the final sample consisted of 54 male and 86 female participants, none of whom was Black. The data from participants who were Black (1 student), who were nonnative speakers of English (2 students), who did not complete the prejudice questionnaire (11 students), or who exceeded the error criterion for the voice key (9 students) were excluded from the analyses.

Procedure

The experimental procedure consisted of two, ostensibly independent, phases. The primary aim of the first phase was to examine the effect of variations in SOAs (permitting controlled processing or not) in a pronunciation task on stereotype activation. The objective of the second phase, presumably an independent task being carried out by a different experimenter, was to assess the respondents' personal stereotypes and level of prejudice.

Phase 1 (pronunciation task). Three types of category primes were used: Black, White, and the letter string, CCC.[2] The target arrays included positive and negative Black stereotypic traits and traits that were not more associated with Blacks than with Whites. Although the primary statistical comparisons focused on the effect of the prime separately for Black stereotypic and nonstereotypic words (and therefore the key hypothesis tests used Black stereotypic and nonstereotypic words as their own controls), special care was taken in the selection of the target traits. Specifically, three lists of synonyms that were associated with the cultural stereotypes of Blacks for both high and low prejudiced people or were nonstereotypic of Blacks were selected on the basis of a series of pilot studies and were equated on word length, word frequency, and favorability.[3]

In the main study, on each trial, participants were presented with an asterisk in the center of the computer

screen for 300 ms (to prepare them for the prime) followed by a blank screen for 500 ms. Next, for participants in the automatic processing condition, the prime appeared for 250 ms followed by a blank screen for 50 ms before the onset of the target word (300 ms SOA). For participants in the controlled processing condition, the prime appeared for 1,950 ms followed by a blank screen for 50 ms before the onset of the target word (2,000 ms SOA). The target array was presented until the voice key was activated. The response latency was recorded and presented on the computer screen for 1,000 ms, followed by a blank screen for 1,000 ms before the next trial.

Five blocks of trials were presented. Each block consisted of 12 trials in which the White, Black, and CCC primes were presented four times. Within a given block, each category was presented with one positive and one negative stereotypical and nonstereotypical trait. Across blocks, each category was presented with five positive and five negative stereotypical and nonstereotypical traits, respectively. The order of presentation of trials within a block and the order of the blocks were randomized for each participant. Over all trials, each category prime was presented multiple (20) times, paired once with each trait (Dovidio, Evans, & Tyler, 1986; Neely, Keefe, & Ross, 1989).

Participants, tested individually, were informed that each trial consisted of an initial asterisk and two words that would appear in sequence. They were instructed to read the first word silently and to read the second target word as quickly and accurately as possible into the microphone. They began with 24 practice trials (using stimuli not included in the main study). Response times were recorded by the computer through the voice-activated microphone. Errors, defined as stutters, mispronunciation of words, reading the prime rather than the target word, or inappropriate triggering of the voice key were recorded by the experimenter who was present throughout this phase of the study. Because the responses of some participants were difficult to monitor (i.e., participants who consistently spoke too softly or exhaled heavily into the microphone before reading the target word), the validity of these responses was questionable. For this reason, the data from participants with nine or more errors (9 participants, who constituted 6% of the sample) were excluded from the analyses.

Phase 2 (prejudice and stereotype assessment). To assess level of prejudice and stereotype attributions, respondents received a questionnaire that included three measures of prejudice and a personal stereotype measure (presented in one of four different orders). To measure prejudice, respondents were asked to complete three prejudice scales: measures of blatant and modern racism and an evaluation thermometer. The blatant racism measure consisted of 20 items adopted from a number of "old-fashioned" Black racism scales (McConahay, Hardee, & Batts, 1981; Pettigrew & Meertens, 1995; Shaw & Wright, 1967). The modern racism measure, which was intended to assess subtle prejudice, consisted of seven items adapted from symbolic and modern racism scales created by Sears (1988) and McConahay et al. (1981). Examples of blatant and modern items, respectively, are "I disapprove of marriage between Blacks and Whites" and "Blacks and other minorities are getting too demanding in their push for equal rights." Both measures of old-fashioned and modern racism used a 7-point Likert-type scale ranging from *strong disagreement* (1) to *strong agreement* (7). A third measure of prejudice, the evaluation thermometer, was also included to examine more global evaluations of the target group (Campbell, 1971; Esses et al., 1993; Judd, Park, Ryan, Brauer, & Kraus, 1995). The evaluation thermometer used a graphic depiction to assess the overall evaluations of groups on a global dimension of favorability and ranges from 0 = *extremely unfavorable* to 100 = *extremely favorable.* Participants were asked to give ratings toward Black and White Canadian men.[4] The secondary purpose of Phase 2 was to examine the participants' personal endorsement of stereotypic traits. To achieve this goal, a measure was included in which participants were asked to estimate the percentage of Black and White Canadian men they believed possessed a specific trait (Brigham, 1971).

Although, as previously described, the selection of target traits was based on extensive pilot testing, the attribution estimate measure also provided an additional check as to whether the target traits used in Phase 1 were appropriately perceived by the current sample as being stereotypical or nonstereotypical. Specifically, *t* tests were conducted on the estimates of each trait for Black and White Canadian men. One trait and its accompanying synonyms from each Stereotypicality × Valence condition were excluded on the basis of the results from the *t* tests (i.e., stereotypic words that did not significantly differentiate between Black and White categories, and nonstereotypic words that did) or the mean percentage of error rate. The list of experimental target words for both the stereotype attribution and stereotype activation analyses appear in the appendix.

RESULTS

Prejudice Measures

Coefficient alpha and mean interitem correlations were computed on the old-fashioned and modern racism scales. Given the acceptable reliabilities (αs = .85 and .72, respectively) and mean interitem correlations (*r*s = .23 and .27, respectively), two indices were created by adding the scores for each participant on the items

for each scale and dividing by the number of items. On the evaluation thermometers, participants' favorability ratings for White Canadian men were divided by their ratings for Black Canadian men. High scores on all indices indicate more prejudice.

To examine the relationship among the three prejudice indices, Pearson product-moment correlations were calculated. The old-fashioned and modern racism scales were significantly related to each other, $r(138) = .63$, $p < .001$. Participants who scored high on the old-fashioned racism scale also scored high on the modern racism scale. The evaluation thermometer was also significantly related to the old-fashioned racism scale, $r(131) = .39$, $p < .001$, and the modern racism scale, $r(131) = .28$, $p < 001$.[5] Participants who rated Black males less favorably than White males scored higher on the racism scales. After each scale score was standardized, a factor analysis with varimax rotation indicated that these three scales loaded on a single factor (eigenvalue = 1.87, accounting for 63% of the variance). Given the high intercorrelations and the results from the factor analysis, a single prejudice index was created by calculating the mean of the *z* scores for each participant on the three scales. Based on a median split, 70 participants (22 men and 48 women) were classified as low and 70 participants (32 men and 38 women) were classified as high in prejudice. Sex of participant and classification as high or low in prejudice were not significantly associated.

Stereotype Activation

On the response latency data set, following the procedures suggested by Ratcliff (1993) and used by Banaji and Hardin (1996) and Blair and Banaji (1996), latencies exceeding three standard deviations plus the mean were identified as outliers and along with responses associated with errors were removed and replaced with missing values. The mean percentages of error and outlier rates were 2.26% and 1.46%, respectively. Remaining response times were subjected to logarithmic transformations. For each participant, the transformed mean response latency for each condition (Positive vs. Negative Traits × Stereotypical vs. Nonstereotypical Traits) were computed for all three primes (Black, White, and CCC). Notwithstanding that all of the analyses were performed on the logarithmic transformed data, the untransformed means are presented in Table 1 and reported in the text.

Although a string array of CCCs was included in the present study and the means are reported in Table 1 for reference, the appropriateness of this type of array as a measure of baseline responding is questionable. Specifically, it must be demonstrated that baseline primes are comparable in all respects, other than their lack of relation to respective social groups to the Black and White primes (see Jonides & Mack, 1984). Mean response latencies, therefore, were used in the present study rather than facilitation scores in which the mean from the Black and White prime conditions is first subtracted from the baseline prime conditions.[6] As with a number of recent studies (Bargh et al., 1996; Blair & Banaji, 1996), we chose to dispense with the facilitation score method because of the difficulties in assessing a true baseline or neutral prime condition (Bargh, Chaiken, Govender, & Pratto, 1992; Fazio, Sanbonmatsu, Powell, & Kardes, 1986). Furthermore, in accordance with earlier studies concerning stereotype attribution and activation, our primary interest was in the comparison of Black versus White primes on stereotype activation (Dovidio et al., 1986; Gaertner & McLaughlin, 1983). This strategy also enabled us to compare directly results related to stereotype activation with results related to stereotype attribution. It should be noted, however, that when facilitation scores were considered, the analysis yielded an identical pattern of means and statistical significance as the present set of analyses because the same value for the CCC condition is subtracted from each condition.

A 2 (prejudice: high vs. low) × 2 (SOA: 300 ms vs. 2000 ms condition) × 2 (prime: Black and White) × 2 (stereotype: Black stereotypic and nonstereotypic target words) × 2 (valence: positive and negative target words) analysis of variance (ANOVA) was performed on the transformed activation scores. Only the prejudice and SOA factors were between subjects; all other factors were within subject. The analysis revealed main effects for SOA, $F(1, 136) = 6.90$, $p < .02$; Stereotype, $F(1, 136) = 5.23$, $p < .03$; and Valence, $F(1, 136) = 13.48$, $p < .001$. In general, participants responded faster under 300 SOA conditions than under 2,000 SOA conditions, *M*s = 495 ms and 520 ms; faster to stereotypes than nonstereotypes, *M*s = 505 ms and 509 ms; and faster to positive than negative traits, *M*s = 505 ms and 510 ms. A Prime × Stereotype interaction was also significant, $F(1, 136) = 5.26$, $p < .03$. Participants' responses to nonstereotypes were faster after viewing White in comparison to Black category primes, *M*s = 506 ms and 512 ms. Participants' responses to Black stereotypes, however, did not differentiate as a function of White or Black category primes, both *M*s = 505 ms.

The above findings, however, are qualified by a Prime × Stereotype × Prejudice interaction, $F(1, 136) = 6.33$, $p < .02$. The means are presented in Table 1. Inconsistent with the set of hypotheses derived from Devine's (1989) position, the effect was not moderated by SOA; the SOA × Stereotype × Prime × Prejudice interaction was not significant, $F(1, 136) = .02$, $p = .89$. As anticipated, for high prejudiced individuals, the Prime × Stereotype interaction was significant, $F(1, 69) = 11.02$, $p < .001$. Consistent

with the hypothesized stereotypic association for high prejudiced persons, participants responded faster to Black stereotypic words following Black primes than White primes, *Ms* = 504 ms versus 511 ms, $F(1, 69) = 4.69$, $p < .04$. High prejudiced participants also responded faster to nonstereotypic words following White primes than Black primes, *Ms* = 508 ms versus 516 ms, $F(1, 69) = 5.33$, $p < .03$.

In contrast to the results indicating stereotypic activation for high prejudiced participants, there was no evidence of differential activation of stereotypic and nonstereotypic words as a function of prime for low prejudiced participants. For low prejudiced participants, the Prime × Stereotype interaction did not approach statistical significance, $F(1, 69) = .02$. Furthermore, consistent with Lepore and Brown's (1997) findings but inconsistent with Devine's (1989) position, there was no evidence of differential stereotype activation for low prejudiced participants across the short and long SOA conditions. The Prime × Stereotype × SOA interaction was not significant when only low prejudiced participants were considered, $F(1, 68) = .32$. In addition, the Prime × Stereotype interaction was not significant within 300 ms SOA conditions, $F(1, 35) = .06$, or 2,000 ms SOA conditions, $F(1, 33) = .48$, separately. Instead, low prejudiced participants responded generally slower following the Black prime than the White prime, *Ms* = 508 ms versus 502 ms, $F(1, 69) = 5.45$, $p < .03$ (see Table 1). These results suggest that high and low prejudiced participants may, in fact, have different representations of Blacks.

TABLE 1: Mean Latencies (in milliseconds) as a Function of Participant's Level of Prejudice, Prime, Trait Stereotypicality, and Processing Conditions (SOA)

	Stereotypes		*Nonstereotypes*	
	300 SOA	*2,000 SOA*	*300 SOA*	*2,000 SOA*
Low prejudice				
Black prime	496	518	498	521
White prime	490	512	491	518
CCC prime	506	515	508	527
High prejudice				
Black prime	490	516	502	528
White prime	497	523	493	521
CCC prime	498	516	503	524

Stereotype Attribution

For each participant, the mean attribution estimates for the stereotypical versus nonstereotypical traits were computed for both Black and White target categories. A 2 (prejudice: high vs. low) × 2 (category: Black and White) × 2 (stereotype: Black stereotypic and nonstereotypic target words) ANOVA, with repeated measures on the last two factors, was performed on these attribution indices. The analysis revealed main effects for Category, $F(1, 138) = 69.67$, $p < .001$, and Stereotype, $F(1, 138) = 39.88$, $p < .001$. In general, participants attributed more traits to Blacks than Whites, *Ms* = 51% and 47%, respectively, and attributed more stereotypes than nonstereotypes, *Ms* = 51% and 48%, respectively. A number of two-way interactions, Prejudice × Category, $F(1, 138) = 8.09$, $p < .01$, Prejudice × Stereotype, $F(1, 138) = 17.03$, $p < .001$, Category × Stereotype, $F(1, 138) = 175.27$, $p < .001$, were also significant.

The above findings, however, were qualified by a significant three-way Prejudice × Category × Stereotype interaction, $F(1, 138) = 20.39$, $p < .001$. The means are presented in Table 2. Simple effects analysis of the high prejudiced participants' responses indicated a highly significant Prime × Stereotype interaction, $F(1, 69) = 146.27$, $p < .001$. As expected, high prejudiced participants attributed stereotypic traits more to Blacks than to Whites, *Ms* = 61% versus 46%, $F(1, 69) = 184.19$, $p < .001$. Moreover, high prejudiced participants attributed nonstereotypic traits somewhat more to Whites than Blacks, *Ms* = 50% versus 46%, $F(1, 69) = 15.80$, $p < .001$. The Prime × Stereotype interaction was also significant, but much less so, for low prejudiced individuals, $F(1, 69) = 41.24$, $p < .001$. These participants also attributed stereotypic traits more to Blacks than Whites, *Ms* = 53% versus 45%, $F(1, 69) = 50.47$, $p < .001$, and nonstereotypic traits more to Whites than Blacks, *Ms* = 49% versus 47%, $F(1, 69) = 4.24$, $p < .05$. Thus, although both high and low prejudiced participants generally endorsed these cultural stereotypic associations, the three-way interaction indicates that high prejudiced participants did so to a greater extent.

Prejudice, Stereotype Activation, and Stereotype Attribution

To examine further the relationship among the prejudice index, response latencies related to stereotype activation on the pronunciation task, and stereotype attributions, Pearson product-moment correlations were calculated. First, however, difference scores were computed by subtracting means in the Black category/stereotype condition from means in the corresponding White category/stereotype condition for both the attribution estimate index and the pronunciation task index. Higher scores on the attribution index indicate more trait attribution to Blacks than to Whites. Higher scores on the pronunciation task indicate slower trait activation related to Blacks than to Whites.

Consistent with the results of the ANOVA in which participants were classified as high or low in prejudice based on the median split, prejudice score (treated as a continuous variable) was significantly correlated with stereotype activation, $r(138) = -.17$, $p < .05$. Participants

TABLE 2: Mean Percentage Trait Attributions as a Function of Participant's Level of Prejudice and Trait Stereotypicality

	Stereotypes	*Nonstereotypes*
Low prejudice		
Blacks	53	47
Whites	45	49
High prejudice		
Blacks	61	46
Whites	46	50

who scored higher on the prejudice scale responded faster to Black stereotypic words after a Black than a White prime. The magnitude of this relationship was comparable for both the 300 SOA condition, $r(67) = -.16$, and the 2,000 SOA condition, $r(69) = -.20$. These findings support the results of the ANOVA that demonstrated that prejudice and stereotype activation were related equivalently across the two SOA conditions.

Furthermore, in accordance with theorizing by Lepore and Brown (1997), individual differences in stereotype attribution was found to be significantly related to stereotype activation, $r(138) = -.16$, $p < .05$. In general, participants who attributed stereotypes to a larger proportion of Blacks relative to Whites responded faster to stereotypes after Black than White primes. Furthermore, this relationship occurred when only responses under 300 SOA conditions were considered, $r(67) = -.25$, $p < .05$. This effect was weaker, but not significantly so, in the 2,000 SOA condition, $r(69) = -.08$.

DISCUSSION

The present research examined the interrelationships among individual differences in racial prejudice, racial stereotyping (i.e., attributions), and stereotype activation. Devine (1989) proposed that because they have been equivalently exposed to and are equally knowledgeable of cultural stereotypes, high and low prejudiced people display equivalent levels of automatic activation of stereotypes. However, under more controlled circumstances, low prejudiced people inhibit manifestations of stereotypes and thus exhibit less stereotyping than do high prejudiced people under these conditions. In contrast, Lepore and Brown (1997) have suggested that when only category priming is involved, because high and low prejudiced people differ in their level of personal endorsement of cultural stereotypes, develop different degrees of associative strength with consensual stereotypes, and potentially have divergent representations, individual differences in prejudice will moderate automatic stereotype activation.

The present study, which used a pronunciation task under automatic and controlled processing conditions, extended this earlier work by not only examining (a) the relationship between prejudice and stereotype activation but also directly examining (b) the relationship between prejudice and stereotype attribution and (c) the relationships among prejudice, stereotype attribution, and stereotype activation. Consistent with the findings of Lepore and Brown (1997), our results revealed that, whereas high prejudiced participants demonstrated Black stereotype activation both under conditions designed to permit conscious control of responses (i.e., 2,000 SOA conditions) and conditions designed to preclude conscious control (i.e., 300 SOA conditions), low prejudiced participants showed no systematic facilitation of Black stereotypic words following a Black prime in either SOA condition.

Beyond conceptually replicating the difference between high and low prejudiced persons in automatic stereotype activation, the present research helps to elucidate the underlying mechanisms by providing direct evidence of the links proposed by Lepore and Brown (1997). Consistent with the research of Devine (1989) and Lepore and Brown (1997), the results of pretesting indicated that high and low prejudiced people were equally knowledgeable about the stereotypes of Blacks that were used in the present study. Furthermore, in accordance with previous research and Lepore and Brown's proposition, high and low prejudiced participants differed in their level of personal endorsement for these stereotypes. Although both groups showed significant stereotypic associations for Blacks, the effect was substantially stronger for high prejudiced participants. The distinction between stereotype knowledge and strength of associative links deriving from frequency of activation of endorsed characteristics lies at the core of the model proposed by Lepore and Brown (1997). Supportive of this model, we found evidence of a relationship between personal stereotype attributions with prejudice and with stereotype activation in general and automatic stereotype activation in particular.

In viewing stereotypes as associations with categories, current network models of stereotyping assume that these links can vary in strength from weak to strong associations (Stephan & Stephan, 1993). It is important to note here the difference between strength of stereotype association and knowledge of cultural stereotypes. Although Devine (1989) found that high and low prejudiced individuals had similar knowledge of cultural stereotypes related to Black categories, it remains possible that these individuals differ in the strength with which they associate stereotypic traits to category representations. Indeed, it is highly plausible that even though people have been similarly socialized in a culture in which stereotypes and negative attitudes toward Blacks may be prevalent, and therefore have acquired the same

knowledge of cultural stereotypes, individual differences exist in their cognitive representations of Blacks.

Two factors that may be related to the strength with which people associate characteristics to a category are the consistency with which they activate the traits upon presentation of the category and the extremity with which they attribute these traits to the category. Both our own theoretical viewpoint and Devine's (1989) results suggest that high prejudiced individuals may express racial stereotypes more frequently and consistently and expose themselves to traditional stereotypes more often. Low prejudiced people, alternatively, may be able to avoid the use of stereotypes personally and by others. Furthermore, research by Fazio and his colleagues (Fazio & Williams, 1986; Powell & Fazio, 1984) has demonstrated the existence of a reliable relationship between attitude extremity and latency of response to attitudes. These findings suggest that more extreme judgments may be more strongly associated with the construct and therefore more accessible. Because high prejudiced people use stereotypes more consistently, engage in repeated activation of stereotypes, and attribute stereotypes more extremely to category members, they may develop associations that are highly accessible and of sufficient strength to produce automatic activation (cf. Bargh et al., 1992; Chaiken & Bargh, 1993). Because low prejudiced people engage in less stereotyping in general and attribute stereotypes less extremely to category members, they may develop weaker associations that are less accessible, or even develop divergent associations (Lepore & Brown, 1997). These individuals are therefore less likely to activate cultural stereotypes automatically.

Although the present results provide convergent evidence for the automaticity of stereotype activation that complements previous research on racial attitudes (Devine, 1989), gender stereotypes (Banaji & Greenwald, 1995; Banaji & Hardin, 1996; Blair & Banaji, 1996), and stereotypes of the elderly (Hense et al., 1995), the findings contribute to a growing inconsistency in the literature about the relationship between attitudes and stereotype activation. Whereas a number of studies have found a significant relationship (Hense et al., 1995; Lepore & Brown, 1997; Wittenbrink et al., 1997), a number of others have not (Banaji & Greenwald, 1995; Devine, 1989; Gaertner & McLaughlin, 1983).

What can account for the discrepant findings concerning the effects of prejudice among the various stereotype activation studies? Why is level of prejudice not consistently found to be related to stereotyping? One relatively straightforward explanation involves the magnitude of the prejudice and stereotyping relationship. In a meta-analysis of 30 hypotheses from 12 different studies of individual differences in racial prejudice and stereotyping, Dovidio et al. (1996) demonstrated a significant but modest relationship ($r = .25$). Further examination of a number of recent studies demonstrates that, although they display a marked variance in the relationship between prejudice and automatic stereotype activation, their effect sizes are also modest. For example, moderate to weak correlations were found in the Wittenbrink et al. (1997) and Lepore and Brown (1997) subliminal category priming studies ($rs = .29$ and $.34$, respectively), Locke et al.'s (1994) study using the Stroop procedure ($r = .33$), the Hense et al. (1995) implicit stereotyping study ($r = .15$), Devine's (1989) study of the implicit hostility attributions ($r = .12$), and the present study of stereotype activation under 300 SOA ($r = .16$). These results suggest that although self-report measures of prejudice have some predictive value in determining who will automatically activate stereotypes, their usefulness is dependent on paradigms that have substantial statistical power or methodological precision. Specifically, procedures may need to be sufficiently sensitive to enable the detection of this generally modest relationship between prejudice and stereotyping.

The pronunciation task used in the present research may represent one such sensitive technique. Unlike most previous studies, this technique permits an examination of stereotype activation using the same type of response (i.e., word pronunciation) and the same procedure while varying a single parameter (SOA) to restrict or allow controlled processing. The effectiveness of the pronunciation paradigm in measuring automatic processing has been demonstrated by a number of different cognitive and social psychologists employing a wide array of stimuli (Bargh et al., 1996; Besner & Chapnik Smith, 1992; De Groot, 1985; Joordens & Besner, 1992; Neely, 1991). Thus, the results from this task, in the 300 ms SOA condition in particular, provide direct evidence of differences in automatic stereotype activation associated with level of prejudice.[7] We assumed that the longer (2,000 ms) SOA condition, in this paradigm as in other response latency paradigms (Neely, 1977, 1991), would reflect controlled processing. Nevertheless, even if one argued that the 2,000 ms SOA condition also represented automatic processing, our basic conclusion would be unchallenged: Automatic stereotype activation is systematically related to individual differences in prejudice and stereotyping.

In conclusion, the present research, besides underlining the importance of methodologically sensitive paradigms, also suggests the importance of exploring variables that may moderate or mediate the relationship between prejudice and automatic stereotype activation. A productive avenue for future research in this area, which converges with other research on attitudes in general (Petty & Krosnick, 1995), might be to under-

stand what characteristics people associate with specific categories, the extremity of those associations, and the strength with which they associate those characteristics with the category. This in turn will permit examination of how these factors moderate accessibility in general, and automatic and controlled stereotype activation specifically. This approach, which reinforces Devine's (1989) distinction between automatic and controlled processes, is consistent with Lepore and Brown's (1997) hypothesis about the role of personal stereotypes, and it further recognizes that for some people, for some characteristics, and for some categories, automatic stereotype activation may not always be inevitable.

APPENDIX
Experimental Target Words

Positive Stereotypes and Synonyms		
1. athlete	sports	muscular
2. musical	rhythm	harmony
3. hip	trendy	cool
4. funny	amusing	witty
5. helpful*	generous*	thoughtful*
Negative Stereotypes and Synonyms		
1. aggressive	tough	hostile
2. poverty	jobless	welfare
3. angry	bitter	mad
4. illegal	criminal	convict
5. uneducated*	ignorant*	illiterate*
Positive Nonstereotypes and Synonyms		
1. cheerful	glad	happy
2. upfront*	outspoken*	frank*
3. elegant	graceful	stylish
4. loyal	faithful	devoted
5. friendly	pleasant	nice
Negative Nonstereotypes and Synonyms		
1. fat*	plump*	chubby*
2. strange	bizarre	weird
3. nasty	cruel	vicious
4. cunning	deceptive	devious
5. nervous	anxious	tense

*Target words not included in the present analyses.

NOTES

1. Although we recognize and acknowledge the preference of many people for the term *African American* or *Black Canadian* instead of *Blacks*, because the present study and much of the related research specifically use the term *Blacks* as a stimulus category, we use that term as well.

2. Although target traits and category primes related to Homosexual and Heterosexual categories were also included, the present article will focus on the results related to Black and White categories.

3. One pilot study (n = 150) identified an initial pool of Black stereotypic and nonstereotypic traits. A second pilot study (n = 20) determined whether these traits were positively or negatively valenced. A third (n = 49) assessed the stereotypicality and favorability of synonyms for the original set of traits. A fourth pilot study (n = 14) confirmed the choice of words representing consensual Black stereotypes and nonstereotypes and demonstrated that perceptions of the cultural stereotypicality of these traits were comparable for high and low prejudiced Whites. Details of these pilot studies are available from the first author.

4. For pragmatic reasons, the focus of the present study was on stereotypes related to Black Canadian men. Because cultural stereotypes related to Black men and women may be divergent, and because general group stereotypes may be associated more strongly with men than with women (Eagly & Kite, 1987), the choice was made to limit the present study to male target persons. However, analyses of stereotype activation including sex of participant resulted in no consistent effects for sex.

5. From a total sample of 140 participants, 7 participants failed to rate the target categories on the evaluation thermometers.

6. The authors would like to thank an anonymous reviewer for suggesting this strategy.

7. In contrast to some of the previous pronunciation studies, the present study used multiple presentations of the same set of priming stimuli. Although a number of recent studies also employ this strategy (e.g., Bargh et al., 1996), it is possible that with this procedure, participants are able to ruminate about the target groups represented by the primes over time. This process then may conceivably produce intentional, rather than automatic, activation of memory contents associated with the categories, which could mitigate differences between the 300 and 2,000 SOA conditions. Support for this possibility would be reflected in a shift in the pattern of responding across trials within the 300 SOA condition: Early trials would reflect automatic activation, whereas late trials would represent a more intentional activation. However, inconsistent with this interpretation, additional analyses demonstrated no such effect for trial block within the present study.

REFERENCES

Allport, G. (1954). *The nature of prejudice.* Reading, MA: Addison-Wesley.

Balota, D., & Chumbly, J. (1984). Are lexical decisions a good measure of lexical access? The role of word frequency in the neglected decision stage. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 340-357.

Balota, D., & Lorch, R. (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.

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

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

Bargh, J., 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., Chaiken, S., Raymond, P., & Hymes, C. (1996). The automatic evaluation effect: Unconditional automatic attitude activation and a pronunciation task. *Journal of Experimental Social Psychology, 32,* 104-128.

Besner, D., & Chapnik Smith, M. (1992). Models of visual word recognition: When obscuring the stimulus yields a clearer view. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 468-482.

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

Brigham, J. (1971). Ethnic stereotypes. *Psychological Bulletin, 76,* 15-38.

Campbell, A. (1971). *White attitudes towards Black people.* Ann Arbor, MI: Institute for Social Research.

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

De Groot, A. (1985). Word-context effects in word naming and lexical decision. *Quarterly Journal of Experimental Psychology, 37,* 281-297.

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

Dovidio, J., Brigham, J., Johnson, B., & Gaertner, S. (1996). Stereotyping, prejudice, and discrimination: Another look. In N. Macrae, C. Stangor, & M. Hewstone (Eds.), *Stereotypes and stereotyping* (pp. 1276-1319). New York: Guilford.

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

Dovidio, J. F., & Fazio, R. H. (1992). New technologies for the direct and indirect assessment of attitudes. In J. Tanur (Ed.), *Questions about survey questions: Meaning, memory, attitudes, and social interaction* (pp. 204-237). New York: Russell Sage.

Eagly, A., & Kite, M. (1987). Are stereotypes of nationalities applied to both women and men? *Journal of Personality and Social Psychology, 53,* 451-462.

Esses, V., Haddock, G., & Zanna, M. (1993). Values, stereotypes, and emotions as determinants of intergroup attitudes. In D. Mackie & D. Hamilton (Eds.), *Affect, cognition and stereotyping: Interactive processes in group perception* (pp. 137-166). San Diego, CA: Academic Press.

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

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

Gaertner, S., & McLaughlin, J. (1983). Racial stereotypes: Associations and ascriptions of positive and negative characteristics. *Social Psychology Quarterly, 46,* 3-30.

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

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

Joordens, S., & Besner, D. (1992). Priming effects that span an intervening unrelated word: Implications for models of memory representation and retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 483-491.

Judd, C., Park, B., Ryan, C., Brauer, M., & Kraus, S. (1995). Stereotypes and ethnocentrism: Diverging interethnic perceptions of African American and White American youth. *Journal of Personality and Social Psychology, 69,* 460-481.

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.

Mackie, D., & Hamilton, D. (1993). Affect, cognition, and stereotyping: Concluding comments. In D. Mackie & D. Hamilton (Eds.), *Affect, cognition and stereotyping: Interactive processes in group perception* (pp. 371-384). New York: Academic Press.

McConahay, J., Hardee, 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.

Messick, D. M., & Mackie, D. M. (1989). Intergroup relations. *Annual Review of Psychology, 40,* 45-81.

Neely, J. (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.

Neely, J. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264-336). Hillsdale, NJ: Lawrence Erlbaum.

Neely, J., Keefe, D., & Ross, K. (1989). Semantic priming in the lexical decision task: Roles of prospective prime-generated expectancies and retrospective semantic matching. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 1003-1019.

Pettigrew, T. F., & Meertens, R. W. (1995). Subtle and blatant prejudice in Western Europe. *European Journal of Social Psychology, 25,* 57-75.

Petty, R., & Krosnick, J. (Eds.). (1995). *Attitude strength: Antecedents and consequences.* Hillsdale, NJ: Lawrence Erlbaum.

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

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

Ratcliff, R., & McKoon, G. (1988). A retrieval theory of priming in memory. *Psychological Review, 95,* 385-408.

Sears, D. (1988). Symbolic racism. In P. Katz & D. Taylor (Eds.), *Eliminating racism: Profiles in controversy* (pp. 53-84). New York: Plenum.

Shaw, M., & Wright, J. (1967). *Scales for the measurement of attitudes.* New York: McGraw-Hill.

Smith, E. (1993). Social identity and social emotions: Toward new conceptualizations of prejudice. In D. Mackie & D. Hamilton (Eds.), *Affect, cognition and stereotyping: Interactive processes in group perception* (pp. 297-316). San Diego, CA: Academic Press.

Stephan, W., & Stephan, C. (1993). Cognition and affect in stereotyping: Parallel interactive networks. In D. Mackie & D. Hamilton (Eds.), *Affect, cognition and stereotyping: Interactive processes in group perception* (pp. 111-136). San Diego, CA: Academic Press.

Wittenbrink, B., Judd, C., & 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.

Received April 23, 1996
Revision accepted September 19, 1997