

Evidence for Racial Prejudice at the Implicit Level and Its Relationship With Questionnaire Measures

Bernd Wittenbrink, Charles M. Judd, and Bernadette Park
University of Colorado

The content of spontaneously activated racial stereotypes among White Americans and the relation of this to more explicit measures of stereotyping and prejudice were investigated. Using a semantic priming paradigm, a prime was presented outside of conscious awareness (BLACK or WHITE), followed by a target stimulus requiring a word–nonword decision. The target stimuli included attributes that varied in valence and stereotypicality for Whites and African Americans. Results showed reliable stereotyping and prejudice effects: Black primes resulted in substantially stronger facilitation to negative than positive stereotypic attributes, whereas White primes facilitated positive more than negative stereotypic traits. The magnitude of this implicit prejudice effect correlated reliably with participants' scores on explicit racial attitude measures, indicating that people's spontaneous stereotypic associations are consistent with their more controlled responses.

Over the past 40 years, opinion surveys have documented substantial changes in racial attitudes among White Americans (Campbell, 1971; Greeley & Sheatsley, 1971; Hyman & Sheatsley, 1956, 1964; Schuman, Steeh, & Bobo, 1985; Taylor, Sheatsley, & Greeley, 1978). On survey questions ranging from the assessment of attitudes toward school desegregation to measures of social distance (e.g., "How would you react if a family member wanted to bring a Black friend to dinner?"), anti-African American sentiment has been in a continuous decline. A parallel trend is also reflected in studies assessing Whites' stereotypes about African Americans. These studies, usually involving White college student samples, show a similar decrease in the number of negative traits stereotypically associated with African Americans (Brigham, 1974; Dovidio & Gaertner, 1986; G. M. Gilbert, 1951; Karlins, Coffman, & Walters, 1969; D. Katz & Braly, 1933).

Likewise, Judd, Park, Ryan, Brauer, and Kraus (1995) reported a series of studies examining intergroup perceptions among African Americans and White Americans. In this work, African American participants consistently showed patterns of responses that are typical of intergroup judgments, demonstrating both out-group homogeneity effects (Park & Judd, 1990; Park & Rothbart, 1982) and ethnocentrism. White Americans' judgments of their in-group were generally consistent with these two phenomena, but White Americans' judgments of their out-

group were systematically at odds with what would be expected from the intergroup literature. Specifically, White Americans tended to say that although they held a stereotype of African Americans, they did not see African Americans as strongly conforming to it (i.e., low perceived stereotypicality), and they reported very positive views of African Americans on tasks such as the estimated prevalence of positive and negative attributes in the group.

The increasingly positive picture that emerges from questionnaire studies such as the Judd et al. (1995) research stands in apparent contrast to results from experimental studies using more indirect measures of Whites' racial attitudes (cf. Donnerstein, Donnerstein, Simon, & Ditrichs, 1972; Duncan, 1976; Gaertner & Dovidio, 1977; McConahay, 1983; Sigall & Page, 1971; Word, Zanna, & Cooper, 1974). In these experiments, White participants reliably manifest discriminatory behavior toward African American targets, for example, evaluating African Americans' job qualifications less favorably or being more likely to judge an African American target's behavior as aggressive. On the basis of a review of these less obtrusive studies, Crosby, Bromley, and Saxe (1980) concluded that "anti-Black sentiments are much more prevalent among White Americans than the survey data lead one to expect" (p. 546). This position was echoed by others (cf. Dovidio & Fazio, 1992; Dovidio & Gaertner, 1986; McConahay, 1986; Pettigrew & Meertens, 1995; Sigall & Page, 1971), and in explanation, researchers have focused on the potential reactivity of common questionnaire measures. Sensitive to societal norms of nondiscrimination and equality, respondents try to avoid appearing prejudiced and adjust their answers on attitude surveys accordingly. More subtle measures of racial attitudes, in contrast, are generally designed to limit such normative pressures.

More recently, a rather different explanation has been suggested, according to which a person may, at the same time, hold positive attitudes toward a social group and nevertheless be influenced by negative group stereotypes (Devine, 1989). This work draws on a distinction commonly made in the cognitive

Bernd Wittenbrink, Charles M. Judd, and Bernadette Park, Department of Psychology, University of Colorado.

This research was partially supported by National Institute of Mental Health Grant R01 MH45049. We thank Gideon Markman for his dedicated assistance with the data collection and Lyle E. Bourne, Jr. and Tony Greenwald for their helpful comments.

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

literature according to which a person's processing of information varies to the extent with which it is *explicit*, a result of intentional control and conscious thought processes, or *implicit*, a result of spontaneous, effortless activation of knowledge contents that is driven by cues in the stimulus environment rather than by an active memory search (cf. Jacoby, Lindsay, & Toth, 1992; Posner, 1978; Roediger, 1990; Schacter, 1987; Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1977; Uleman & Bargh, 1989). Following this distinction, it is argued that stereotypic knowledge may affect social judgment and behavior differently, depending on whether the response is based on a controlled, conscious consideration of available information or whether stereotypic knowledge is activated spontaneously, outside of the perceiver's conscious control. Specifically, discrepancies between controlled responses to a given social group and implicit influences of stereotypic knowledge may arise if stereotypic associations that come to mind spontaneously differ from one's explicit attitudes.

To date, most of the work concerned with this distinction has focused on demonstrating that stereotypic knowledge in general, and racial stereotypes in particular, may in fact be activated effortlessly and influence subsequent judgments unbeknownst to the perceiver (Banaji & Greenwald, 1995; Banaji, Hardin, & Rothman, 1993; Devine, 1989; Fazio, Sanbonmatsu, Powell, & Kardes, 1986; Gaertner & McLaughlin, 1983; D. T. Gilbert & Hixon, 1991; Macrae, Stangor, & Milne, 1994; Perdue & Gurtman, 1990). Much less work exists that directly assesses the relationship between implicit and explicit measures of stereotyping and prejudice.

Perhaps the best known study that has examined this relationship is Devine's (1989) work, in which she argued that all White Americans know of and automatically activate the culturally shared negative stereotype of African Americans. People low in prejudice must then invoke their own personal (more positive) beliefs regarding African Americans by means of a controlled process. The implications of this argument are that there are no individual differences in the associations that are automatically activated by the category *Blacks* and, therefore, that the degree of prejudice evidenced on an implicit, or automatic, task will be uncorrelated with the degree of prejudice exhibited on an explicit, or controlled task. To test this "dissociation" argument, Devine first showed that participants high and low in prejudice (determined by scores on the Modern Racism Scale [MRS]; McConahay, Hardee, & Batts, 1981) differed in the relative numbers of positive to negative qualities they listed as characteristic of African Americans. That is, the MRS predicted differences in explicitly generated lists of valenced attributes.

Next, in Devine's (1989) study, participants were subliminally exposed to words associated with the target group, African Americans, such as *busing*, *oppression*, *slavery*, *jazz*, and *basketball*. Some participants saw a large number of these words, and some only a few. Then participants read about and interpreted the behavior of an ambiguously hostile individual, whose ethnicity was not specified. If a large number of African American-related words had been seen in the subliminal priming task, the ambiguous target was judged to be more hostile than if a small number of such words were seen. Importantly, the extent to which this was true did not depend on the participants' scores

on the MRS. Thus, Devine (1989) concluded that implicit stereotyping had occurred and that it had occurred to the same degree for all participants, regardless of their level of explicit prejudice.

We have two primary concerns with this study. First, although in the subliminal task, Devine (1989) was careful to avoid words that were direct semantic associates of *hostile*, many of the words presented were at least indirectly related to the concept "hostility." *Busing*, *oppression*, *nigger*, *ghetto*, *slavery*, and *prejudice* are certainly linked in memory to knowledge of hostility. To the extent that this task activated the general concept of hostility (not as it relates specifically to the stereotypic qualities of African Americans, but simply as a general concept, primed by, for example, *busing* and *oppression*), then it is reasonable to expect both that the ambiguous target would be seen as hostile and that this would not depend on MRS scores. Hostility had been primed for all participants (as a general construct), thereby masking any individual differences in the strength of association of hostility to the group African Americans. A second problem is that the ethnicity of the ambiguous target was left unspecified, and thus, the results are arguably a demonstration of implicit stereotyping, given that the target was not specifically identified as a member of the group (and the default assumption would presumably be that the target was White).

Thus, a fair evaluation of the dissociation hypothesis proposed by Devine (1989) would require a task that more directly measures the degree of association of stereotypic attributes with the category African Americans in an implicit manner and that tests whether this is related to explicitly measured stereotypes. In fact, not only would such a task be useful to examine whether controlled responses are indeed uncorrelated with what is activated at the automatic level, but it would at the same time allow us to explore the issue of strategic manipulation of the responses in the Judd et al. (1995) studies.

A different procedure to get at spontaneous activation of stereotypic knowledge has been developed by Dovidio, Gaertner, and their colleagues (Dovidio, Evans, & Tyler, 1986; Gaertner & McLaughlin, 1983). This task is a variation on Meyer and Schvaneveldt's (1971) classic procedure demonstrating semantic priming effects. In the Meyer and Schvaneveldt procedure, participants first see a single word, the prime (e.g., BREAD) and then are presented with a letter string, the target (e.g., BUTTER), to which they have to respond with, for example, a word-nonword judgment. Response latencies are facilitated by semantic associations between the prime and the target stimulus. A common explanation for this by now well established finding (for a review, see Neely, 1991) derives from the concept of spreading activation. Presumably, activation of the prime spreads to semantically related concepts and thus reduces the time required for the activation of related targets to reach recognition threshold (cf. Neely, 1977; Posner & Snyder, 1975).

Dovidio et al. (1986) used similar reasoning when asking their participants about what attributes they associated with White Americans and African Americans. Using a semantic priming task in which participants had to decide whether a target attribute "could ever be true" or "was always false" of the primed category, Dovidio and his colleagues showed that White American participants were slower to indicate that positive attributes could be true of the Black prime than the White prime and faster to indicate that negative attributes could be true of

the Black prime. Thus, their reaction time results suggest that positive attributes showed weaker, and negative attributes stronger, associations with African Americans than with Whites.

Although the response latency paradigm offers an appealing approach to measuring the spontaneous components of stereotype knowledge, the particular paradigm used by Dovidio et al. (1986) does not provide for a clear differentiation between such spontaneous associations and knowledge that is retrieved deliberately. This is because, in their procedure, both priming and target stimuli are clearly visible, and the judgment task itself requires participants to focus on the relation between various trait attributes and the social groups African Americans and Whites. Thus, participants in these experiments were generally aware that the study "had something to do with prejudice" (Gaertner & McLaughlin, 1983, p. 27), and it is therefore reasonable to assume that participants did reflect on their explicit racial beliefs. In fact, Dovidio et al. recognized their procedure's limitations in assessing solely automatic, spontaneous responses, stating that "it was likely that 'controlled processes' . . . were involved to some extent" (Dovidio et al., 1986, p. 35) in participants' decisions. They argued, however, that managed responses were unlikely to be operating because they did in fact find evidence of prejudice. More recently, however, Judd et al. (1995) used this same method and failed in three separate studies to replicate the prejudice findings of Dovidio et al. Given these more recent results and the nature of the task, it is unlikely that it assesses solely implicit effects of stereotypic knowledge.

In response to these concerns with previous research, we designed an experimental task with two primary goals in mind. First, we wanted to provide a more accurate assessment of Whites' spontaneous, unintentionally activated associations for the group of African Americans. Second, we wished to determine how these spontaneous knowledge aspects relate to people's explicit racial attitudes. Our goal of limiting the influence of controlled thought processes on participants' judgments required us to use an experimental task in which participants would not have to make decisions explicitly related to the target group of African Americans. In fact, we intended to measure participants' associations with the target group without them ever consciously encountering any reference to either African or White Americans. To accomplish this goal, we used a semantic priming task in which participants were asked to make word-nonword judgments to various trait attributes that were preceded by either a subliminally presented group prime (BLACK or WHITE) or a subliminally presented neutral prime. Participants were unaware that the social groups, African Americans and Whites, were even relevant to the experimental task. Of primary interest was the speed with which the word-nonword judgment could be made for positive and negative trait attributes stereotypic of the groups African Americans and Whites (and in each case counterstereotypic of the other group), relative to a neutral prime condition. That is, we examined facilitation in the lexical-decision judgments when a group prime preceded the trait word, as opposed to the neutral prime, and compared the magnitude of this facilitation as a function of whether the trait word was (a) stereotypic of either African Americans or White Americans and (b) positive versus negative in valence (see illustration in Figure 1).

Following Judd and Park (1993), we wished to separately

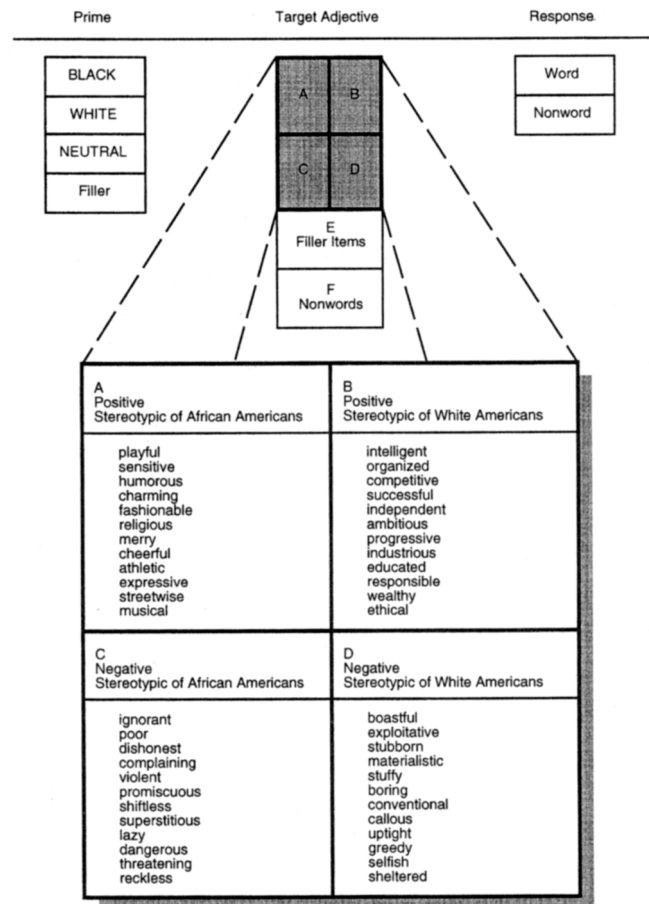


Figure 1. The lexical-decision task consisted of word-nonword judgments to a target item. Besides fillers (E) and nonwords (F), the stimulus pool for this target item included trait adjectives that were either stereotypic of African Americans (and counterstereotypic of Whites; A and C) or stereotypic of Whites (and counterstereotypic of African Americans; B and D). Within each set of stereotypic items, half of the items were positively valenced (A and B), and half were negative in valence (C and D). The target item was preceded by a subliminally presented prime that either referred to one of the two social groups (African Americans and Whites), was neutral, or was a filler.

assess stereotyping versus prejudice effects, and our design, therefore, crossed stereotypicality of the target items with their valence. The inclusion of all four types of items (positive stereotypic, positive counterstereotypic, negative stereotypic, and negative counterstereotypic; Cells A–D in Figure 1) enabled us to test for stereotyping effects that are orthogonal to prejudice effects and vice versa. That is, we wished to assess *implicit stereotyping*, by examining whether items stereotypic of a given group were more strongly associated at an implicit level with that group than an alternative group, and *implicit prejudice*, by examining whether negative attributes were more strongly associated than positive ones at an implicit level for the African American relative to White target group. Although this approach works in theory, in practice it is complicated by the fact that the African American stereotype as it exists in our culture is largely negative. Looking at the attributes in Figure 1 (those

used in this research), one could argue that the negative African American items (Cell C), and the positive White American items (Cell B) are really what constitute the stereotype of these two groups in our society and that the other two types of items are more weakly associated with the groups. In other words, prejudice in this case may be reflected in the very attributes that Whites consider to be stereotypic of African Americans.

In the literature, prejudice is widely defined as a negative attitude toward an out-group (Ashmore, 1970; Dovidio & Gaertner, 1986; Esses, Haddock, & Zanna, 1993; Stephan, 1985; Stroebe & Insko, 1989). Like attitudes more generally (Eagly & Chaiken, 1993), prejudice can be seen as having affective, cognitive, and behavioral components. The *affective* component consists of negative feelings or emotions toward the group; the *cognitive* component consists of negatively valenced beliefs about the group; the *behavioral* component refers to behavioral experiences with members of the group. Although these components tend to be correlated with each other, they certainly need not be entirely redundant. Thus assessments of the cognitive component of prejudice may or may not yield results that perfectly map on to behavioral or affective individual differences.

Because we were interested in the stereotypic associations that participants implicitly make to African Americans, and because we believed that those stereotypic associations were likely to depend on the valence of attributes in question, our task necessarily focused on the cognitive component of prejudice. That is, the procedure that we developed permits us to assess the stereotypic associations that participants have stored in long-term memory and then to ask the further question of whether those stereotypic associations are negatively valenced in the case of the out-group and positively valenced in the case of the in-group, thus revealing prejudice (at least the cognitive component of prejudice). This approach to prejudice assessment is entirely consistent with others who have focused on the more cognitive aspect of prejudice. Stroebe and Insko (1989), for example, offered the following summary for the relationship between stereotypes and prejudice, as it has been conceptualized in the literature:

How are stereotypes toward some group related to prejudice? Information processing approaches as well as consistency theories assume that the attitude toward an attitude object is related to the attributes perceived as associated with that object and the positive or negative evaluation of these attributes. For example, if an American perceives the Germans as "scientifically minded," "industrious" and "intelligent" (stereotype), and if he evaluates these traits positively, he is likely to hold a positive attitude towards the Germans. (Stroebe & Insko, 1989, p. 9)

Our procedure would seem less appropriate if our goal was to examine more generalized affective responses that participants have to African Americans, because the semantic priming paradigm that we use necessarily examines the semantic associations that have been stored in long-term memory. Nevertheless, it is possible in our procedure to ask not only whether the out-group prime (i.e., BLACK) facilitates the negatively valenced attributes that are stereotypic of that group but also, more broadly, whether we get facilitation of all negatively valenced attributes, regardless of whether they are stereotypically associated with African Americans.

The important point, however, is to recognize that these are two alternative definitions of prejudice, focusing on different aspects of an attitude. Our task seems particularly appropriate if one focuses on the cognitive component, asking about the stereotypic associations that participants implicitly make to African Americans and then asking the further question of whether these are negatively valenced. Although it is possible in our task to ask about the more generalized or affective component, we suspected that implicit prejudice would be revealed only on the attributes that are stereotypic of the group primed.

In light of this discussion, we expected to find implicit stereotyping on the part of our participants, so that there was facilitation of African American stereotypic items following the BLACK prime and of White American stereotypic items following the WHITE prime. Additionally, however, we expected that the pattern of this stereotypic facilitation would depend on the valence of the item, thus revealing implicit prejudice. That is, implicit prejudice would be revealed if facilitation on stereotypic African American items following the BLACK prime was particularly found for negatively valenced items and if facilitation on stereotypic White American items following the WHITE prime was particularly found for positively valenced items.

Method

Overview

White American participants were told that they would complete three unrelated experimental tasks. First, they completed a name-identification task, in which they were asked to identify the ethnic group membership of individuals on the basis of their first names. The purpose of this task was to strengthen the association between the ethnic labels that would subsequently be used as group primes and the relevant ethnic groups.

After this task, participants completed a reaction time procedure, in which they were asked whether various target sequences of letters on a computer screen constituted a word or a nonword (see Figure 1). Each trial on this lexical-decision task (LDT) was preceded by a prime that referred to one ethnic group or the other (BLACK or WHITE), a foil prime (e.g., TABLE), or a neutral nonword prime (e.g., XXXXX). These primes were presented for 15 ms and were immediately followed by a mask for a 2-s interval before the lexical-decision trial. Target stimuli during the LDT were either adjectives that were stereotypic of White Americans (and counterstereotypic of African Americans), adjectives that were stereotypic of African Americans (and counterstereotypic of White Americans), irrelevant attributes, or nonwords. Of the stereotypic and counterstereotypic attributes, half were positively valenced, and half were negatively valenced. All stereotypic and counterstereotypic target words occurred with both of the two group primes and with the neutral prime. Differences in response latencies were examined as a function of whether the stereotypic and counterstereotypic target attributes followed a group prime versus a neutral prime. These latency differences were taken as evidence of the associative strength between group labels and target attributes, which in turn would reflect the likelihood with which a given attribute would be activated spontaneously by a group reference. Finally, participants completed a questionnaire that included a variety of explicit questionnaire measures of stereotyping and prejudice.

Participants

Eighty-eight participants were recruited for this study from the introductory psychology participant pool at the University of Colorado. They

received course credit for their participation. African American students were excluded from the sample.

Pretests

To conceal the actual purpose of the LDT, we intended to present the priming stimuli outside of participants' conscious awareness. As is commonly done to obtain such a "subliminal" stimulus exposure, we planned to present the priming stimuli for a brief period of time, followed immediately with a visual mask (xxxxx). Results from previous research indicate that the effectiveness of subliminal stimuli in semantic priming is increased when the stimuli remain detectable, albeit unidentifiable (cf. Holender, 1986). Two pretests were conducted, to decide on the appropriate prime presentation duration. The first pretest was informal and exploratory. We presented primes at various exposure intervals, and we queried participants both about whether they saw a prime flash on the computer screen before the lexical-decision trial and about whether they were able to identify it if they had seen a flash. The purpose of this pretest was to permit us to make an informed choice. Once the choice was made, a second pretest was conducted to verify that primes presented for the chosen duration could not in fact be identified by participants.

To ensure accurate timing of the stimulus exposure on the computer screens, the choice of alternative prime durations was constrained by the time required to refresh the screen and reposition the display beam. With a screen-refresh rate of 66.67 Hz for the monitors used in the experiment, the cycle duration was 15 ms, and we therefore included 15-ms, 30-ms, and 45-ms intervals in the first pretest. The results of this pretest suggested that participants both detected and identified the prime routinely at 45 ms. At 30 ms, they were consistently able to detect the prime, but prime identification was only slightly above chance. At 15 ms, detection was more difficult, albeit somewhat above chance, but correct identification of the prime virtually never occurred. On the basis of these informal results, we made a preliminary choice of 15 ms for the prime duration and then undertook a more formal second pretest to verify that the primes at this presentation duration could not be identified.¹

Twelve participants were recruited for the second pretest. They first completed the name-identification task, just as participants would in the main study. Each participant then completed 150 trials from the LDT, using the same primes and target letter sequences that would be used in the main study, with the prime duration at 15 ms. Unlike the instructions that we actually gave to participants in the main study, these pretest participants were told that words would briefly be flashed on the screen before the masking sequence of Xs and that they should attempt to identify them. On one sixth of the trials, the prime was WHITE; on one sixth of the trials the prime was BLACK; filler primes were used on the remaining two thirds of the trials. Each LDT trial in this pretest was followed by a prime-identification query. In total then, across the 12 participants, there were 1,800 trials with prime-identification queries given on each trial. Of these, 300 identification queries followed each of the two group primes. In response, these participants identified the prime as the word *black* 9 times. They identified the prime as the word *white* 9 times as well. In each case, 8 of these 9 identifications were correct identifications. Thus, out of 300 trials, correct prime identification for each of the two group primes occurred 8 times out of 300, or on 2.67% of the trials. Our interpretation of these results is that even when explicitly told that words would be flashed and asked to identify those words, participants were basically unable to do so at this prime-presentation duration. In the main study, participants were neither informed about the prime nor asked to identify it, making it all the more unlikely that they could or would spontaneously identify the primes.

Procedure

Upon their arrival in the laboratory, participants were told that they would be participating in a number of unrelated studies. First they would

participate in a study of judgmental accuracy, in which they would be asked to identify the ethnic group memberships of different individuals, on the basis of their first names. They were also told, before beginning this first study, that they would participate in a second study, investigating "how people visually recognize words when reading a text." They were further told that in this second study, they would be asked to judge, as quickly and as accurately as possible, whether sequences of letters appearing on a computer screen constituted a word or not. By emphasizing the ostensibly different purpose of the second task, we led participants to believe that the two tasks were entirely unrelated. Finally, they were told that the third task involved completing a questionnaire.

The first task was included so that participants would associate the lexical references *black* and *white* with the relevant ethnic groups, when these labels were presented as primes during the lexical-decision trials. The inclusion of this task circumvented the use of more explicit instructions for participants to "think of African Americans when seeing 'Black' and White Americans when seeing 'White,'" which would have jeopardized our goal of concealing the experiment's actual purpose. Participants were given 20 first names on a questionnaire and were asked, for each one, to identify whether the individual was Black or White. Half of the names were stereotypically African American names (e.g., Lamont or Keisha), and half were stereotypically White American names (e.g., Mark or Elizabeth). For each name, participants indicated their answer by checking either Black or White.

After this task, participants were seated in front of a computer and informed that they would be asked to make judgments about a large number of letter sequences that appeared on the computer screen. They were told that on each trial, they would first see a string of Xs. This string would then be followed by a letter sequence, and they should judge whether the sequence constituted a word or not, making their judgment as quickly and accurately as possible. As described below, participants then were shown 242 lexical-decision trials.

After these trials, participants were asked to complete a questionnaire that included a variety of explicit measures of ethnocentrism and stereotyping. These are described below. At the conclusion of this session, participants were fully debriefed and dismissed.

LDT Stimuli

Presentation of experimental stimuli and data collection was controlled by the PSYSCOPE software package (Cohen, MacWhinney, Flatt, & Provost, 1993) on Apple Macintosh Quadra 605 computers equipped with 14-in. (35-cm) Apple color monitors.

During the experimental administration of this task, participants were seated at a distance of approximately 50 cm from the computer screen and were asked, on each trial, to focus on a fixation point (+), which was presented in the center of the computer screen in 18-point Times Macintosh font. (The same was true of all other stimuli, which were always presented in uppercase letters.) The fixation point appeared for 1,000 ms and was immediately followed by the prime. After 15 ms, the prime was overwritten by a masking stimulus (xxxxx), which remained on the screen for 2,000 ms.² The masking stimulus was then substituted

¹ In addition to the main study reported here, a second study was conducted, setting the prime duration at 30 ms. The results of this study are presented briefly in the General Discussion section.

² Results from previous research indicate that the effectiveness of "subliminal" stimuli in semantic priming is increased when the stimulus onset asynchrony (SOA) between prime and target stimuli is long rather than short (Balota, 1983; Fowler, Wolford, Slade, & Tassinari, 1981). On the basis of these findings, we decided to use a relatively long SOA of 2,000 ms. Unlike situations in which the primes are supraliminal, the length of the SOA should have no effect on the probability that controlled processes are induced.

Table 1
Explicit Prejudice Measures: Representative Items and Coefficient Alpha Reliabilities

Representative item	Alpha
Modern Racism Scale (McConahay et al., 1981)	.802
"Blacks are getting too demanding in their push for equal rights."	
"Over the past few years, the government and news media have shown more respect to Blacks than they deserve."	
Pro-Black scale (I. Katz & Hass, 1988)	.771
"Black people do not have the same employment opportunities that Whites do."	
"Too many Black people still lose out on jobs and promotions because of their skin color."	
Anti-Black scale (I. Katz & Hass, 1988)	.859
"On the whole, Black people don't stress education and training."	
"The root cause of most of the social and economic ills of Blacks is the weakness and instability of the Black family."	
Diversity Scale	.672
"There is a real danger that too much emphasis on cultural diversity will tear the United States apart."	
"The desire of many ethnic minorities to maintain their cultural traditions impedes the achievement of racial equality."	
Discrimination Scale	
"Blacks are ultimately responsible for the state of race relations in this country."	
"More and more, Blacks use accusations of racism for their own advantage."	.885

by one of the target letter sequences, and participants were required to indicate whether the target stimulus formed a correct word. The target sequence was erased from the screen after 250 ms, with the computer pausing until the participant had responded by pressing one of two keys, labeled *yes* and *no*. Participants' two index fingers had previously been positioned over the two response keys.

The experiment included four different sorts of primes. On one fourth of the trials, the prime was the word *white*. For one fourth of the trials, the prime was the word *black*. One fourth of the trials used a neutral nonword, identical to the masking stimulus (xxxxx). Finally, on one fourth of the trials, filler nouns that made no reference to person categories (e.g., *table*, *lemon*, or *summer*) were used as the primes. Each prime (i.e., BLACK, WHITE, neutral, or filler item) occurred on 58 different trials. Of these, on 10 trials, the target sequence of letters following the prime constituted a nonword. The remaining 48 trials for each prime type were followed by targets that were, in fact, words. For the first three types of primes (WHITE, BLACK, and neutral), the target words were 48 attributes that were either stereotypic of White Americans or stereotypic of African Americans. Additionally, half of the target words in each stereotypicality set were positively valenced, and half were negatively valenced. These target words, identified by their stereotypicality and their valence, are given in the lower portion of Figure 1. They were chosen on the basis of data reported by Judd et al. (1995) and of an extensive review provided by Dovidio and Gaertner (1986) of studies using D. Katz and Braly's (1933) adjective checklist approach for the assessment of ethnic stereotypes. This set of 48 target words was fully crossed with all three prime types. For the 58 trials that involved filler primes, 10 were followed by nonwords, and the remaining 48 trials involved positively and negatively valenced target adjectives that could not be used to refer to persons (e.g., *sunny*, *polluted*, or *deserted*). The full set of four prime types by 58 trials (232 trials) was presented in an order that was individually randomized for each participant. Additionally, 10 practice trials were presented initially, involving both filler and neutral primes. No interval separated these 10 practice trials from the 232 experimental trials.

Explicit Questionnaire Measures

After the LDT, participants were asked to complete a questionnaire containing five different explicit measures of ethnic stereotyping and prejudice. Measures included the MRS (McConahay et al., 1981), the Pro-Black and Anti-Black subscales (I. Katz & Hass, 1988), and two scales developed specifically for this study, one tapping participants'

beliefs about the value of ethnic diversity in our society (Diversity Scale) and one tapping beliefs about patterns of discrimination in our society (Discrimination Scale). Table 1 contains representative items from each of these scales, as well as internal-consistency reliability estimates.

Results

We first present the analysis of the latencies from the LDT. In this analysis, we define and test various contrasts, to examine implicit stereotyping and prejudice, as revealed by response facilitation differences. Subsequently, we discuss relationships between the implicit results and the explicit questionnaire measures of prejudice. To do this, we present correlations between the various within-subject contrasts that have been defined using the LDT data and the explicit questionnaire results.

Analysis of LDT Response Latencies

Response latencies are naturally characterized by positive skew and the prevalence of outliers (Ratcliff, 1993). Accordingly, response latencies from the LDT faster than 150 ms and slower than 1,500 ms were recoded as missing values. As a result, 173 individual latencies were recoded as missing values (1.09%). Additionally, to correct the positive skew that naturally characterizes response latency data, both log and inverse or reciprocal transformations of the data were explored. Following the former transformation, the data still showed evidence of considerable positive skew. Hence, an inverse transformation was conducted on all latencies before analyses. The means we report in the text and tables have been retransformed, however, back into the millisecond metric.³

Response latencies to items following a group prime (BLACK or WHITE) were analyzed by comparing them to latencies to the identical items following the neutral prime (xxxxx). In this way, differences in speed of word recognition as a function

³ Analyses using the log transformation yielded results that were very similar to those that we report under the inverse transformation.

of the individual items were eliminated. Thus, the dependent variables in all analyses were response facilitation differences, subtracting the response latency for an item following one or the other group prime from the latency for that same item following the neutral prime. Greater facilitation is indicated by more positive differences. The facilitation differences were analyzed as a function of group prime, item stereotypicality (African American vs. White American items), and item valence, with all factors varying within participants. The primary analysis we report collapsed across the 12 individual items within each of the stereotypicality by valence cells of the design.

The mean facilitation scores are presented in Table 2. The analysis of these data yielded a reliable Prime \times Item Stereotypicality interaction, $F(1, 87) = 4.57, p = .035$, a marginally reliable Item Stereotypicality \times Item Valence interaction, $F(1, 87) = 2.93, p = .090$, and a reliable Prime \times Item Valence interaction, $F(1, 87) = 4.03, p = .048$. The first of these effects represents what we might call implicit stereotyping. Specifically, facilitation was reliably greater when an item's stereotypicality matched the prime that preceded it (mean facilitation = 7.96) than when it was preceded by a prime from the other group (mean facilitation = 2.16). The marginally reliable Item Stereotypicality \times Item Valence interaction tells us that positive items that were stereotypic of White Americans and negative items that were stereotypic of African Americans showed stronger facilitation than positive African American items and negative White American items. Because this interaction does not include group prime, the most parsimonious interpretation of the interaction is that the positive White American and negative African American items are simply more strongly associated with both person categories. Finally, the Prime \times Item Valence interaction tells us that positively valenced items showed larger facilitation following the WHITE prime than the BLACK prime, whereas negatively valenced items showed larger facilitation following the BLACK prime.

In interpreting these results, it is important to remember that the facilitation means reported in Table 2 represent comparisons of response latencies following the group primes relative to trials on which the exact same items were preceded by the neutral prime. More positive values indicate that the items were recognized as words more quickly following one or the other group prime than following the neutral prime. Accordingly, it seems important to determine whether these mean values indicate facilitation by the relevant group primes in an absolute sense, by testing whether the individual cell means in Table 2 differ reliably from zero. These tests reveal that group primes facilitated responses reliably only in those cells in which a

positively valenced White American item was preceded by a WHITE prime, $F(1, 87) = 6.02, p = .016$, and in which a negatively valenced African American item was preceded by a BLACK prime, $F(1, 87) = 6.80, p = .011$.

Accordingly, the reliable implicit stereotyping effect that we reported (i.e., the Prime \times Item Stereotypicality interaction) seems largely attributable to the strength of the facilitation for the positive White American items following the WHITE prime and of the negative African American items following the BLACK prime. Thus, as we suggested in the introduction, stereotypes of African Americans and White Americans in our society tend to have rather different valences associated with them. Our facilitation results confirm this speculation.

The overall analyses of these facilitation data are equivalent to tests of various single-degree-of-freedom within-subject contrasts. The Prime \times Item Stereotypicality interaction that was found to be reliable, which we labeled *implicit stereotyping*, involves the contrast weights presented in the top panel of Table 3. We can define other contrasts that are also of interest, both from the point of view of analyzing the LDT data further and from the point of view of correlations that will subsequently be reported with the explicit prejudice measures.

The contrast weights in the second panel of Table 3 capture the theoretical definition of *implicit prejudice*, as defined in the introduction. Namely, with this contrast, we are examining whether the strength of the stereotypic associations with the two group primes depends on the valence of the item. In other words, given that we are only examining items that are stereotypic of the group that has been primed, is the resulting facilitation dependent on valence, so that it is stronger for the positively valenced White American items, following the WHITE prime, and the negatively valenced African American items, following the BLACK prime? This contrast proved to be reliable in these data, $F(1, 87) = 4.81, p = .031$. Like the implicit stereotyping effect, note that the reliability of this implicit prejudice contrast is largely driven by the large facilitation associated with the upper left and lower right cells of the design, namely, when positively valenced White American items follow the WHITE prime and when negatively valenced African American items follow the BLACK prime. Given the centrality of these cells to both the implicit stereotyping and implicit prejudice contrasts, it is clear that a strong distinction cannot be made between the two in the present data. The evaluative nature of the ethnic stereotypes we are examining results in an empirical confound between stereotyping and prejudice, even though in theory (Judd & Park, 1993), the two are distinguishable.

In addition to an overall test of the implicit prejudice contrast weights presented in the second panel of Table 3, we also tested the two separate components of out-group derogation and in-group favoritism that combine to form the implicit prejudice contrast. The out-group derogation component examined only the facilitation difference due to valence when African American items followed the BLACK prime. The in-group favoritism component examined the facilitation difference due to valence when the White American items followed the WHITE prime. Out-group derogation with these facilitation data proved to be reliable, $F(1, 87) = 3.78, p = .055$, in-group favoritism did not, $F(1, 87) = 2.55, p = .114$.

Table 2
Mean Response Facilitation (in Milliseconds) Following a Group Prime, Relative to a Neutral Prime

Item valence	Items stereotypic of African Americans		Items stereotypic of White Americans	
	BLACK prime	WHITE prime	BLACK prime	WHITE prime
Positive	-0.72	-2.72	5.17	14.47
Negative	17.74	3.94	2.26	0.36

Table 3
Contrast Weights Used in Lexical-Decision Task Analysis

Item valence	Items stereotypic of African Americans		Items stereotypic of White Americans	
	BLACK prime	WHITE prime	BLACK prime	WHITE prime
I. Implicit stereotyping contrast ^a				
Positive	1	-1	-1	1
Negative	1	-1	-1	1
II. Implicit prejudice contrast				
Positive	-1	0	0	1
Negative	1	0	0	-1
IIa. Out-group derogation				
Positive	-1	0	0	0
Negative	1	0	0	0
IIb. In-group favoritism				
Positive	0	0	0	1
Negative	0	0	0	-1
III. Alternative stereotyping contrast				
Positive	0	0	-1	1
Negative	1	-1	0	0
IV. Generalized prejudice contrast ^b				
Positive	-1	1	-1	1
Negative	1	-1	1	-1

^a Prime \times Item Stereotypicality interaction. ^b Prime \times Valence interaction.

The third panel of contrast weights in Table 3 represents an alternative conceptualization of implicit stereotyping.⁴ Given that the positively valenced White American items are more stereotypic of White Americans than are the negatively valenced ones and given that the negatively valenced African American items are more strongly stereotypic of African Americans than are the positively valenced ones, then perhaps an alternative way of examining implicit stereotyping is to ask whether these items are more strongly facilitated by their appropriate group primes than by the primes from the other group. Given the pattern of means and the strong facilitation in the upper right and lower left cells of the design, it is not surprising that the contrast that tests this interaction is reliable, $F(1, 87) = 6.66, p = .012$. What will be of perhaps greater interest is to examine whether it shows different patterns of correlations with the explicit questionnaire measures than our original implicit stereotyping contrast that ignores item valence.

The final contrast, represented by the weights in the fourth panel of Table 3, that we tested examined whether prejudice generalized across item stereotypicality.⁵ That is, here we are asking whether the BLACK prime facilitates all negatively valenced items more than positively valenced ones, regardless of whether the items are included in the African American stereotype. Similarly, we are asking whether the WHITE prime facilitates all positively valenced items regardless of their stereotypicality. Actually, we previously reported the test of this contrast, because it is identical to the Prime \times Item Valence interaction, $F(1, 87) = 4.03, p = .048$.⁶

Although we have given different names to the four sets of contrast weights (and their components) contained in Table 3, they are, to varying extents, partially redundant, as a result of the fact that stereotyping and prejudice are partially redundant in the present context. That is, the fact that each of these con-

trasts is reliable is largely due to the substantial facilitation of the positively valenced White American items following the WHITE prime and of the negatively valenced African American items following the BLACK prime. To describe the redundancy of these contrasts, participant scores on each of them were computed by weighting the individual facilitation scores for each participant by the contrast weights. These individual contrast scores were then correlated with each other. The resulting correlation matrix is presented in Table 4.

Relationships Between LDT Results and Explicit Questionnaire Measures

As a final task, participants completed a questionnaire that contained the various explicit prejudice scales presented in Table 1. To examine the relationships between the explicit measures and the implicit stereotyping and prejudice results from the LDT

⁴ We thank Myron Rothbart for suggesting the conceptualization of stereotyping represented by this contrast.

⁵ We thank David Hamilton for suggesting that we examine the definition of prejudice represented by this contrast.

⁶ All of the analyses that we have reported collapsed across attribute items, within levels of valence and stereotypicality, and treated participant as the random factor in the design. Analyses were also conducted in which attribute item was treated as a random factor, collapsing across participants. In these analyses, attribute valence and stereotypicality were between-item factors, and prime varied within attributes. The results were largely consistent with those we have reported. Specifically, both the implicit stereotyping (Prime \times Item Stereotypicality interaction) and the implicit prejudice contrasts were reliable. Thus, the results we report seem generalizable not only across participants but also across attribute items.

Table 4
Correlations Among Lexical-Decision Task Contrast Scores

Contrast	I	II	IIa	IIb	III	IV
I. Implicit stereotyping	—					
II. Implicit prejudice	.34**	—				
IIa. Out-group derogation	.43***	.86***	—			
IIb. In-group favoritism	.08	.76***	.48***	—		
III. Alternative stereotyping	.82***	.63***	.62***	.38***	—	
IV. Generalized prejudice	.29**	.68***	.57***	.55***	.79***	—

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

data, we correlated the explicit scores with the various within-subject contrast scores computed from the LDT data. The resulting correlation matrix is presented in Table 5, with the top half of the table presenting the intercorrelations among the explicit measures and the bottom half containing the correlations between the explicit measures and the LDT contrasts. The various LDT contrasts are defined by the weights used in Table 3 and the labels used in Tables 3 and 4. Higher scores on these various within-subject contrasts indicate that participants are displaying more stereotypic and more prejudicial implicit responses to the priming stimuli. With the exception of the Pro-Black subscale from I. Katz and Hass (1988), all of the explicit questionnaire measures are coded so that higher scores indicate greater prejudice.

A number of things deserve comment about these correlations. First, all of the explicit questionnaire measures intercorrelated substantially and in the expected direction.

Second, and most dramatically, with the exception of the Anti-Black subscale from I. Katz and Hass (1988), all of the explicit questionnaire measures were reliably related to all of the implicit contrast scores that involved item-valence differences and that, therefore, could be thought of as assessing implicit prejudice in various ways. Thus, the MRS, the Pro-Black subscale, and our own Diversity and Discrimination scales were reliably related in every case with the implicit prejudice contrast (II), the out-group derogation component of this contrast (IIa), the in-group

favoritism component of this contrast (IIb), and the more generalized prejudice contrast (IV).

Third, the implicit prejudice contrast (II) correlated somewhat more highly with the explicit questionnaire measures than did the generalized prejudice contrast (IV), suggesting that the preferred definition of prejudice on this task concerned valence differences in the content of the group-specific stereotype that is activated. This is consistent with the conceptualization we offered in the introduction.

Fourth, both of the components of the implicit prejudice contrast were reliably related to the explicit measures, although the out-group derogation contrast was, across the board, somewhat more highly related to the questionnaire measures than was the in-group favoritism contrast. Importantly, however, the full implicit prejudice contrast correlated slightly more highly with the explicit measures than either of its components did alone, suggesting that prejudice as assessed by the questionnaire measures took into account not only participants' implicit evaluations of the out-group but also their implicit evaluations of the in-group.

Fifth, the two LDT contrasts that do not involve valence differences (i.e., what we labeled the *implicit stereotyping* contrast and the *alternative stereotyping* contrast) were not, with one exception, reliably related to the explicit questionnaire measures. Thus, implicit prejudice, as picked up on our LDT with subliminal group primes, seems to importantly involve the dif-

Table 5
Correlations Among Explicit and Implicit Measures

Measure	1	2	3	4	5
Explicit measure					
1. Modern racism	—				
2. Pro-Black	-.56***	—			
3. Anti-Black	.52***	-.30**	—		
4. Diversity	.56***	-.46***	.49***	—	
5. Discrimination	.73***	-.67***	.61***	.58***	—
Implicit contrast scores					
I. Implicit stereotyping	.13	.03	-.05	.03	-.08
II. Implicit prejudice	.41***	-.33**	.17	.32**	.32**
IIa. Out-group derogation	.40***	-.32**	.12	.28**	.28**
IIb. In-group favoritism	.25*	-.21*	.16	.23*	.23*
III. Alternative stereotyping	.22*	-.12	.06	.15	.10
IV. Generalized prejudice	.24*	-.23*	.15	.21*	.25*

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

ferences in response facilitation to the two primes as a function of the valence of item. The one exception is that the alternative stereotyping contrast does correlate with the MRS, although more weakly than the contrasts that involve item valence. Additionally, implicit prejudice depends crucially on the stereotypicality of an item, because the correlations weakened when we moved from the implicit prejudice contrast (II) to the generalized prejudice contrast (IV), which includes the valence difference on the items that are not stereotypic of the primed group.

One final analysis seemed appropriate. Given the centrality and reliability of the facilitation for the positively valenced White American items following the WHITE prime and for the negatively valenced African American items following the BLACK prime, one might wonder whether differing degrees of facilitation for just these cells of the design might also correlate substantially with the explicit measures. After all, these cells of the LDT design certainly play a crucial role in all of the contrasts that proved reliable and that were correlated with the explicit measures. Perhaps all of the meaningful variation in these implicit results occurs simply in these two cells.

To examine this possibility, we correlated facilitation in these two cells with the explicit measures. Because these facilitation scores in fact represent differences in response latencies between neutral prime trials and group prime trials, they are within-subject contrast scores similar to the contrast scores we used to this point. Interestingly, this contrast, involving facilitation only in the upper right and lower left cells of the design as laid out in Table 3, correlated less well with the explicit questionnaire measures than did the full implicit prejudice contrast (II). For instance, the correlation of the implicit prejudice contrast with the MRS was .41. The correlation of the facilitation contrast for just the upper right and lower left cells of the design with the MRS was .31. Although this difference certainly is not significant, it does suggest that although facilitation in the cells that are given negative contrast weights in the implicit prejudice contrast (i.e., positively valenced African American items following the BLACK prime and negatively valenced White American items following the WHITE prime) was not reliable, nevertheless, meaningful variation existed in the latency differences of these cells that predicted the explicit measures. For instance, the correlation between facilitation for the positively valenced African American items following the BLACK prime and scores on the MRS was $-.32$ ($p < .01$). Participants who displayed lower levels of prejudice on the MRS were more likely to show response facilitation when a positively valenced African American item followed the BLACK prime, compared with when it followed the neutral prime, even though in an absolute sense, there was no facilitation on average in this cell of the design.

General Discussion

Our goal in this study was to develop a completely unobtrusive measure of White American participants' associations with the social categories of African Americans and White Americans. We wanted to do so both to be able to assess the content of those implicit and spontaneous associations and to relate them to questionnaire measures of prejudice. These goals were motivated both by prior work suggesting the reduction or even absence of prejudice as measured by explicit questionnaires

assessing White Americans' attitudes toward African Americans (e.g., Judd et al., 1995) and by conclusions reached in the literature concerning the dissociation between implicit responses and explicit questionnaire results (e.g., Devine, 1989).

The implicit task that we used focused on the speed with which participants could identify a letter string as a word or nonword, depending on a preceding prime that was presented outside of conscious awareness. The letter strings that we were interested in were words that were stereotypic of either White Americans or African Americans and that were either positively or negatively valenced. Evidence of implicit stereotyping and prejudice was revealed by facilitation in the identification of these words when they followed one or the other of the group primes relative to when they followed the neutral prime.

The results from this task revealed evidence of implicit stereotyping and prejudice. Specifically, item identification was significantly facilitated when positively valenced White American items followed the WHITE prime and when negatively valenced African American items followed the BLACK prime. The reliable facilitation in these cells of the design rendered a number of contrasts that examined differences in facilitation between these cells and other cells reliable as well. For instance, we found reliably greater facilitation when an item's stereotypicality matched the prime that preceded it than when it was preceded by a prime from the other group (i.e., the implicit stereotyping contrast). Additionally, among items that were stereotypic of the group primed, valence was found to make a difference: Positively valenced White American items were facilitated more than were negatively valenced ones following the WHITE prime, and negatively valenced African American items were facilitated more than were positively valenced ones following the BLACK prime (i.e., the implicit prejudice contrast).

In combination, these results suggest that the task we developed was relatively successful at eliciting spontaneous stereotypic reactions to the social categories of White Americans and African Americans and that these implicit responses show evidence of stereotyping and prejudice.⁷ Additionally, and in keeping with our speculations in the introduction, the stereotypic contents of these associations are colored by their valences, so that stereotyping and prejudice on the implicit level are conceptually intertwined.

Part of our motivation in testing the various contrasts from the implicit task, over and above the tests of absolute facilitation in the specific cells of the design, was to measure participants'

⁷ Two aspects of our design may have facilitated these implicit effects. First, the initial naming task, in which participants were asked to identify the ethnicity of various first names, may have made the ethnic categories, and their associations, more salient than would have been the case if this task had not preceded the lexical-decision trials. However, the presence of this task should make those associations, whatever their content, simply more accessible. It should not influence the content of what is implicitly activated. Nor can it explain the differences in facilitation due to the prime and valence manipulations. Second, because of the repeated trials aspect of our design, each target item followed three primes, the two group primes and the neutral prime. This means that earlier trials with the same item quite likely facilitated later trials for that item. Because order was randomized at the level of the individual, however, this facilitation due to order is independent of the prime facilitation that we have shown.

implicit stereotyping and prejudice in various ways and correlate these with the explicit questionnaire measures. Individual participant scores were computed on each of the individual within-subject contrasts that were tested, and these were correlated with the explicit results. These correlations strongly suggested that the implicit associations revealed by our task were consistent with participants' explicit questionnaire responses. Whenever the implicit contrasts took into account the valence of the implicit associations with the two primed groups, the correlations with all of the explicit measures, with the exception of the Anti-Black scale from I. Katz and Hass (1988), were significant. Focusing on differences among the explicit measures, the correlations with the implicit prejudice results were consistently most strong with the MRS (McConahay et al., 1981). Focusing on differences among the various implicit contrasts, the one that we labeled *implicit prejudice* consistently showed the highest correlations. These correlations were slightly higher than either the out-group derogation or the in-group favoritism component contrasts were, suggesting that the explicit measures were affected by both sorts of implicit responses. Additionally, correlations involving the full implicit prejudice contrast were somewhat larger than those that only included facilitation in the lower left and upper right cells of the design, suggesting that meaningful individual variation exists in the other cells of the design, at least in terms of the association of that variation with the explicit measures, even though facilitation was not found in an absolute sense in these cells.

Our confidence in these results (both the demonstration of implicit prejudice and its correlation with explicit questionnaire responses) results not only from their consistency and statistical reliability but also from the fact that we conducted an additional study in which they were replicated.⁸ Actually, this additional study was conducted prior to the one reported above. The implicit task that we used in the additional study was identical to that of the present experiment, with one major exception: The prime duration was 30 ms rather than 15 ms. At this longer prime duration, there was some evidence that prime identification was above chance. Because of this, we made the decision to rerun the study, shortening the prime duration to 15 ms, so that we could more confidently claim that prime identification virtually never occurred. The results in this other study, using the longer prime duration, were virtually identical to those in the current study. Specifically, we found reliable implicit stereotyping and prejudice and the relevant contrast scores from these implicit results were once again reliably and consistently related with the explicit questionnaire measures. The results even replicated the findings that the MRS showed the largest correlations with the implicit contrasts and that the implicit prejudice contrast showed the largest correlations with the explicit measures.

There is one additional reason for mentioning this replication study. After participants had completed the lexical-decision trials but before they completed the prejudice questionnaire, they were seated at the computer and taken through the task developed by Dovidio et al. (1986) and used in the Judd et al. (1995) studies. Recall that this task asked participants whether positively and negatively valenced attributes could ever be true of the primed groups, Blacks and Whites, with the primed category presented supraliminally. In keeping with the results found in three studies in Judd et al. (1995), no evidence was found

for prejudice among our participants on this task. Additionally, the within-subject contrast that measures prejudice on the task was uncorrelated with either the explicit questionnaire measures or the implicit prejudice contrasts from the LDT. Thus, consistent with our speculations in the introduction, we do not believe that the Dovidio et al. task measures responses at the implicit level. Perhaps somewhat surprisingly, it also seems unrelated with responses revealed by the explicit questionnaires.

Given the data presented here and elsewhere, what can we conclude about the nature of the relationship between implicit and explicit measures of prejudice? If we wanted to measure someone's "true" level of prejudice, must we rely on implicit measures, or given our results, can explicit questionnaire measures serve the same purpose? We know that the explicit measures used in this study, namely, attitude questionnaires such as the MRS, are related to implicit prejudice, and therefore, scores on these explicit scales could be defensibly used as an indicator of participants' underlying or implicit attitudes. Note, however, that although this type of explicit measure can tell us about the ordering of participants relative to one another, it cannot reveal an individual's level of prejudice in an absolute sense. That is, these racial attitude scales examine participants' beliefs about particular social issues, such as the fairness of affirmative action or busing, and are able to scale individuals with respect to how positively or negatively their attitudes are toward African Americans. But the scales do not assess how respondents evaluate African Americans, either absolutely or relative to how they evaluate White Americans. Thus, although these attitude questionnaires are meaningful for assessing variation among participants in their level of prejudice, they cannot reveal in any absolute sense the presence or absence of prejudice.

To determine whether prejudice exists in terms of mean responses to explicit measures, one needs, for example, judgments of the perceived prevalence of positive and negative attributes among African Americans and White Americans. Such measures were used by Judd et al. (1995). They consistently found that White American college students reported very positive stereotypes of African Americans (more positive than those given to the White American target group and nearly as positive as those given by African American participants). Judd et al. argued against a simple social desirability or management interpretation of these data. Instead they suggested that the responses reflected patterns of socialization so that the White American students are taught to attempt to treat others fairly and to avoid using ethnicity or race in making judgments about others. These participants' judgments were influenced, it was argued, by their ideological beliefs about fair treatment of others and, in all likelihood, by their images of themselves as egalitarian and fair individuals. Note, however, that at the level of the mean judgments, the current LDT data reveal that White American participants possess a much more negative view of African Americans at the implicit level than that revealed at the explicit level in the judgment tasks used by Judd et al.

So the pattern of results across studies is that absolute levels of prejudice are revealed in our implicit task, but these disappear when participants complete an explicit task that permits the

⁸ A full write-up of this study is available from the authors.

assessment of whether they demonstrate prejudice in any absolute sense. At the same time, other explicit measures of stereotyping and prejudice that do not permit one to determine absolute levels of prejudice (i.e., the MRS) are relatively highly correlated with the implicit measure. Given this consistency at the explicit and implicit levels, then, an important question concerns those factors that are responsible for the mean shift in the unprejudiced direction found on those explicit measures that permit an absolute assessment of prejudice.

It may be that all participants are motivated to portray themselves as more egalitarian and unprejudiced than their implicit reactions suggest. This then would result in a positive mean shift with preservation of rank ordering of participants between the explicit and implicit tasks. On the other hand, the correlation between the two is certainly not perfect, so it is entirely reasonable that the mean shift depends on the relative importance of participants' ideological beliefs concerning equality and fairness. Thus, both participants' implicit beliefs as well as their ideological beliefs about the role of race and ethnicity in our society contribute to their explicit questionnaire responses.

In the introduction, we argued that in Devine's (1989) work, individual differences in the strength of association of hostility to the group of African Americans may have been masked because the hostility concept was primed for all participants through related associates. We therefore felt it is questionable whether the implicit task used in her research specifically assessed beliefs about African Americans in a manner analogous to the explicit task. The present findings are important at a theoretical level because they suggest that the model offered by Devine is at least partly implausible. Devine argued that the (largely negative) cultural stereotype of African Americans is automatically activated when the category itself is activated. At this automatic (or implicit) level, all individuals show evidence of the same negative associations with African Americans. According to Devine, low-prejudice participants go on to activate a personal, more positive stereotype, in a controlled fashion, when asked to give explicit race-relevant responses.

Our data indicate that participants do differ in the strength of negative versus positive associations with African Americans, relative to White Americans, even at an implicit level and that this variation is predicted by scores on the MRS, the same explicit measure used in Devine's (1989) research. High- and low-prejudice participants are activating different associations, even at the implicit level. The mean shift we observed in overall prejudice from the implicit task (reported here) to the explicit trait judgment task (as in Judd et al., 1995) is consistent with Devine's notion that a more negative stereotype acquired from socialization in our culture is automatically activated and must then be altered with conscious activation of a more positive personal stereotype. But we disagree that there is no variation in the automatic component of what is activated. Our data indicate it is strongly correlated with explicit racial attitude questionnaires.

In sum, our data indicate the presence of implicit stereotyping and implicit prejudice and indicate that the level of implicit prejudice is well predicted by explicit questionnaire measures of racial attitudes. The data argue for the importance of assessing the same construct at the implicit and explicit level and for the importance of a truly implicit measure. They argue

against different information coming to mind as a function of the nature of the task (implicit vs. explicit) for those who score high versus low in prejudice. We are particularly interested in the way in which participants' knowledge at the implicit level, beliefs regarding the group and the current interethnic social situation, and ideological beliefs regarding equality are integrated into explicit judgments of what the group is like, as well as judgments about specific individual group members. The current studies, in conjunction with other recent research, make it clear that White Americans' views of African Americans are multifaceted. Knowledge and beliefs at both the implicit and explicit level, as well as ideological and societal beliefs, all influence such judgments. With future research, we will hopefully move closer to a more complete and accurate characterization of this perception process.

References

- Ashmore, R. D. (1970). Prejudice: Causes and cures. In B. E. Collins (Ed.), *Social psychology: Social influence, attitude change, group processes, and prejudice* (pp. 245-339). Reading, MA: Addison-Wesley.
- Balota, D. A. (1983). Automatic semantic activation and episodic memory encoding. *Journal of Verbal Learning and Verbal Behavior*, 22, 88-104.
- 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., & Rothman, A. J. (1993). Implicit stereotyping in person judgment. *Journal of Personality and Social Psychology*, 65, 272-281.
- Brigham, J. C. (1974). Views of Black and White children concerning the distribution of personality characteristics. *Journal of Personality*, 42, 144-158.
- Campbell, A. (1971). *White attitudes toward Black people*. Ann Arbor, MI: Institute for Social Research.
- 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.
- Crosby, F., Bromley, S., & Saxe, L. (1980). Recent unobtrusive studies of Black-and-White discrimination and prejudice: A literature review. *Psychological Bulletin*, 87, 546-563.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, 56, 5-18.
- Donnerstein, E., Donnerstein, M., Simon, S., & Ditrachs, R. (1972). Variables in interracial aggression: Anonymity, expected retaliation, and a riot. *Journal of Personality and Social Psychology*, 22, 236-245.
- 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., & Fazio, R. H. (1992). New technologies for the direct and indirect assessment of attitudes. In J. M. Tanur (Ed.), *Questions about questions* (pp. 204-236). New York: Russell Sage Foundation.
- 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.
- Duncan, B. L. (1976). Differential social perception and attribution of intergroup violence: Testing the lower limits of stereotyping of Blacks. *Journal of Personality and Social Psychology*, 34, 590-598.

- Eagly, A. H., & Chaiken, S. (1993). *The psychology of attitudes*. Fort Worth, TX: Harcourt Brace Jovanovich.
- Esses, V. M., Haddock, G., & Zanna, M. P. (1993). Values, stereotypes, and emotions as determinants of intergroup attitudes. In D. M. Mackie & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping: Interactive processes in group perception* (pp. 137-166). San Diego, CA: Academic Press.
- Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology*, 50, 229-238.
- Fowler, C. A., Wolford, G., Slade, R., & Tassinari, L. (1981). Lexical access with and without awareness. *Journal of Experimental Psychology: General*, 110, 341-362.
- Gaertner, S. L., & Dovidio, J. F. (1977). The subtlety of white racism, arousal, and helping behavior. *Journal of Personality and Social Psychology*, 35, 691-707.
- Gaertner, S. L., & McLaughlin, J. P. (1983). Racial stereotypes: Associations and ascriptions of positive and negative characteristics. *Social Psychology Quarterly*, 46, 23-30.
- 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.
- Gilbert, G. M. (1951). Stereotype persistence and change among college students. *Journal of Abnormal and Social Psychology*, 46, 245-254.
- Greeley, A. M., & Sheatsley, P. B. (1971). Attitudes toward racial integration. *Scientific American*, 225, 13-19.
- Holender, D. (1986). Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: A survey and appraisal. *The Brain and Behavioral Sciences*, 9, 1-23.
- Hyman, H. H., & Sheatsley, P. B. (1956). Attitudes toward desegregation. *Scientific American*, 195, 35-39.
- Hyman, H. H., & Sheatsley, P. B. (1964). Attitudes toward desegregation. *Scientific American*, 211, 16-23.
- Jacoby, L. L., Lindsay, D., & Toth, J. P. (1992). Unconscious influences revealed: Attention, awareness, and control. *American Psychologist*, 47, 802-809.
- Judd, C. M., & Park, B. (1993). Definition and assessment of accuracy in social stereotypes. *Psychological Review*, 100, 109-128.
- Judd, C. M., Park, B., Ryan, C. S., Brauer, M., & Kraus, S. (1995). Stereotypes and ethnocentrism: Interethnic perceptions of African American and White American college samples. *Journal of Personality and Social Psychology*, 69, 460-481.
- Karlins, M., Coffman, T. L., & Walters, G. (1969). On the fading of social stereotypes: Studies in three generations of college students. *Journal of Personality and Social Psychology*, 13, 1-16.
- Katz, D., & Braly, K. W. (1933). Racial stereotypes of one hundred college students. *Journal of Abnormal and Social Psychology*, 28, 280-290.
- 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.
- 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. (1983). Modern racism and modern discrimination: The effects of race, racial attitudes, and context on simulated hiring decisions. *Personality and Social Psychology Bulletin*, 9, 551-558.
- McConahay, J. B. (1986). Modern racism, ambivalence, and the Modern Racism Scale. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 91-125). Orlando, FL: Academic Press.
- 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.
- Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, 90, 227-234.
- 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.
- Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264-336). Hillsdale, NJ: Erlbaum.
- Park, B., & Judd, C. M. (1990). Measures and models of perceived group variability. *Journal of Personality and Social Psychology*, 59, 173-191.
- Park, B., & Rothbart, M. (1982). Perception of outgroup homogeneity and levels of social categorization: Memory for the subordinate attributes of ingroup and outgroup members. *Journal of Personality and Social Psychology*, 42, 1051-1068.
- Perdue, C. W., & Gurtman, M. B. (1990). Evidence for the automaticity of ageism. *Journal of Experimental Social Psychology*, 26, 199-216.
- Pettigrew, T. F., & Meertens, R. W. (1995). Subtle and blatant prejudice in Western Europe. *European Journal of Social Psychology*, 25, 57-75.
- Posner, M. I. (1978). *Chronometric explorations of mind*. Hillsdale, NJ: Erlbaum.
- Posner, M. I., & Snyder, C. R. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola Symposium* (pp. 55-85). Hillsdale, NJ: Erlbaum.
- Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin*, 114, 510-532.
- Roediger, H. L. (1990). Implicit memory: Retention without remembering. *American Psychologist*, 45, 1043-1056.
- Schacter, D. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 501-518.
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, 84, 1-66.
- Schuman, H., Steeh, C., & Bobo, L. (1985). *Racial attitudes in America*. Cambridge, MA: Harvard University Press.
- 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.
- Sigall, H., & Page, R. (1971). Current stereotypes: A little fading, a little faking. *Journal of Personality and Social Psychology*, 18, 247-255.
- Stephan, W. G. (1985). Intergroup relations. In G. Lindzey & E. Aronson (Eds.), *Handbook of social psychology* (3rd ed., Vol. 2, pp. 599-658). New York: Random House.
- Stroebe, W., & Insko, C. A. (1989). Stereotypes, prejudice, and discrimination: Changing conceptions in theory and research. In D. Bar-Tal, C. F. Grauman, A. W. Kurganski, & W. Stroebe (Eds.), *Stereotypes and prejudice: Changing conceptions* (pp. 3-34). New York: Springer.
- Taylor, D. G., Sheatsley, P. B., & Greeley, A. M. (1978). Attitudes toward racial integration. *Scientific American*, 238, 42-49.
- Uleman, J. S., & Bargh, J. A. (Eds.). (1989). *Unintended thought*. New York: Guilford Press.
- Word, C. O., Zanna, M. P., & Cooper, J. (1974). The nonverbal mediation of self-fulfilling prophecies in interracial interaction. *Journal of Experimental Social Psychology*, 10, 109-120.

Received May 25, 1995

Revision received April 24, 1996

Accepted May 28, 1996 ■