

BLACK AMERICANS' IMPLICIT RACIAL ASSOCIATIONS AND THEIR IMPLICATIONS FOR INTERGROUP JUDGMENT

Leslie Ashburn–Nardo
University of Kentucky

Megan L. Knowles
Northwestern University

Margo J. Monteith
University of Kentucky

Many—though not all—black Americans exhibit an implicit evaluative preference for whites relative to blacks (e.g., Livingston, 2002). Are such biases meaningfully related to blacks' explicit attitudes and actual intergroup judgments? In the present study, 83 black participants who believed they would complete an intellectually challenging task with a partner rated their preferences for (fictitious) black and white potential partners. The less strongly participants implicitly preferred their ingroup, the lower their preference for a black vs. a white work partner. The magnitude of this relationship held even when controlling for explicit attitudes that were related to partner preference. Implicit biases were associated with explicit attitudes regarding black, but not white, persons and with system–justifying ideology (Jost & Banaji, 1994).

Order of authorship was determined alphabetically. Leslie Ashburn–Nardo is now in the Department of Psychology at Indiana University – Purdue University Indianapolis.

This manuscript is based on an undergraduate honors thesis submitted by Megan L. Knowles to the University of Kentucky. Portions of this research were presented at the Annual Convention of the Society for Personality and Social Psychology, Savannah, Georgia, February 2002. The research was supported by a grant from the Society for the Psychological Study of Social Issues Grants–in–Aid program. We would like to thank Julietta Rice for her assistance with data collection and Russ Fazio for his comments on an earlier version of this manuscript.

Address correspondence to: Leslie Ashburn–Nardo, Indiana University – Purdue University Indianapolis, Department of Psychology, 402 North Blackford Street (LD 124), Indianapolis, IN 46202–3275.

Decades of research have documented the negativity that many white Americans associate with black Americans. Many whites harbor and consciously endorse prejudiced attitudes about blacks (for recent data, see Schuman, Steeh, Bobo, & Krysan, 1997). Additionally, whites who have consciously adopted egalitarian and nonprejudiced belief systems often continue to show evidence of activation and application of stereotypes (e.g., Bargh, Chen, & Burrows, 1996; Devine, 1989; Lepore & Brown, 1997) and more negative implicit attitudes in relation to blacks than to whites (e.g., Fazio, Jackson, Dunton, & Williams, 1995; Greenwald, McGhee, & Schwartz, 1998; Wittenbrink, Judd, & Park, 1997).

Far fewer studies have examined blacks' own attitudes toward and evaluative associations regarding their own race. Perhaps the lack of empirical attention to this issue, relative to the attention devoted to understanding whites' racial biases, is because researchers have assumed that blacks, like most other social groups, would be inclined to favor and to evaluate positively their racial ingroup. Indeed, reviews of the literature demonstrate that ingroup preferences are remarkably robust (e.g., Mullen, Brown, & Smith, 1992), and that such biases emerge even when assignment to one's ingroup has just been made on an arbitrary basis (e.g., Tajfel, Billig, Bundy, & Flament, 1971). This strong tendency to favor the ingroup after being arbitrarily assigned to it is even evident on implicit measures (Ashburn–Nardo, Voils, & Monteith, 2001). This suggests that people are mentally prepared to evaluate more positively their ingroups relative to outgroups. One might assume that this bias would be especially likely when the outgroup of interest is one that has historically discriminated against the ingroup.

Studies that have examined blacks' collective self-esteem and racial identity have in fact revealed that blacks appear to hold their race in high esteem, and they appear to be able to separate their personal ingroup views from how they believe society views their race (e.g., Crocker, Luhtanen, Blaine, & Broadnax, 1994; Rowley, Sellers, Chavous, & Smith, 1998). Those studies, however, relied on survey methods and explicit (i.e., consciously controlled) measures. A different picture has begun to emerge with the few studies that have examined blacks' implicit racial associations. In an unpublished dissertation, Spicer (1999) found considerable variability in blacks' implicit racial preferences though, overall, participants showed a significant preference for whites over blacks. The same was true among black participants in an Internet-based study conducted by Nosek, Banaji, and Greenwald (2002).

Livingston (2002) recently reported two studies that examined the variability in blacks' performance on two implicit measures of attitudes toward blacks versus whites. Participants in both samples reported very strong levels of explicit ingroup bias, but significant evidence of implicit bias was not obtained on either implicit measure employed. Nonetheless, there was considerable variability in implicit bias scores, and these were related systematically to participants' reports of the extent to which they believed others undervalue their group (Study 1a) and are prejudiced toward it (Study 1b). In short, participants who perceived greater negativity toward their group were more prone to favoring the outgroup (whites) over the ingroup (blacks). In addition, outgroup contact was related to implicit biases, and evidence that perceived negativity mediates the relation between outgroup contact and implicit biases was obtained. In other words, participants who reported less contact with whites were more prone to perceiving negativity toward blacks and, in turn, to favoring whites over blacks implicitly.

Although it was certainly not the case that all blacks exhibited more positive associations with whites than with blacks, the fact that, as a group, blacks exhibited far less ingroup bias than whites is consistent with an environmental associationist view (Devine, 1989; Karpinski & Hilton, 2001). Specifically, because of the prevalence of positive societal images associated with whites and negative societal images associated with blacks, some consistency across whites' and blacks' implicit associations is to be expected.

Many blacks' implicit biases are also consistent with system justification theory (Jost & Banaji, 1994). Jost and Banaji defined system justification as the "psychological process by which existing social arrangements are legitimized, even at the expense of personal and group interest" (p. 2). In other words, people—whether they are members of low- or high-status groups—attempt to justify the social hierarchy observed within their society. As a result, both low- and high-status group members may exhibit a preference for the high-status group relative to the low-status group, and members of both groups may endorse negative stereotypes regarding the low-status group. Jost and Banaji (1994) argued that these preferences are especially likely to emerge with measures that tap into implicit associations. Recent findings, consistent with system justification theory, have revealed less implicit ingroup bias among members of low- versus high-status groups (Jost, Pelham, & Carvalho, 2002; Rudman, Feinberg, & Fairchild, 2002) and, more

germane to the present study, among blacks versus whites (Nosek et al., 2002).

In the current research, we sought to add to the underdeveloped literature concerning blacks' implicit racial biases in two ways. First, we examined the nature of these associations and their relation with various explicit attitudinal measures. Second, and more importantly, we investigated whether blacks' scores on an implicit measure simply reflect dormant associations or whether they have greater meaning—such as in predicting preferences for a black versus a white partner in an intellectually challenging task. We know from the literature on whites' racial biases that what is reflected in implicit associations need not translate into beliefs (e.g., Devine, 1989), even if such biases are manifested in judgments of target group members (e.g., Dovidio, Kawakami, Johnson, Johnson, & Howard, 1997; Fazio et al., 1995). Is the same true of blacks, or are their implicit biases divorced from their actual intergroup judgments?

MEASURING IMPLICIT ASSOCIATIONS

There are numerous techniques for measuring implicit associations in relation to race, from priming procedures (e.g., Fazio et al., 1995; Gaertner & McLaughlin, 1983; Payne, 2001; Perdue, Dovidio, Gurtman, & Tyler, 1990; Wittenbrink et al., 1997) to considering how word fragments are completed (Gilbert & Hixon, 1991) to considering whether one uses abstract versus concrete words to describe an action (e.g., von Hippel, Sekaquaptewa, & Vargas, 1997) to considering the way and the degree to which one explains stereotype-congruent versus incongruent behavior (Sekaquaptewa, Espinoza, Thompson, Vargas, & von Hippel, 2003) to physiologically based procedures (Vanman, Paul, Ito, & Miller, 1997) (for a review, see Fazio & Olson, 2003). What these measures have in common is their ability to limit consciously controlled responding, although no measure is entirely process-pure (i.e., measures are subject to both controlled and automatic influences) (Jacoby, 1991). The Implicit Association Test (IAT; Greenwald et al., 1998), however, has been shown to be rather impervious to response manipulation attempts that are under conscious control (Banse, Seise, & Zerbes, 2001; Egloff & Schmukle, 2002; Kim & Greenwald, 1998). Because most of the researchers who have attempted to learn about blacks' implicit racial attitudes have used

the IAT, we decided to use this technique as well in order to facilitate comparison with the extant literature on the topic.¹

We also decided to use the IAT because it provides a relative measure of preference for whites versus blacks. This is advantageous because, for example, one might have favorable implicit associations for both whites and blacks on independent measures of those associations yet exhibit a preference for one group over the other on a relative measure (cf. Nosek & Banaji, 2001). Because our main interest was in determining whether implicit associations were related to preference for a white versus black interaction partner, an implicit measure that also taps into relative preference was desirable.

How does the IAT provide a relative index? The IAT is a computer-based dual-categorization task in which participants categorize exemplars (e.g., Brandon, Tyrone) of two social categories (e.g., white, black) as well as exemplars (e.g., sunshine, stink) of two evaluative categories (e.g., pleasant, unpleasant). On critical trials the social and evaluative categories are combined and participants' response latencies are recorded, yielding a measure of the strength of association between the social and evaluative categories. If individuals respond faster when ingroup exemplars are paired with pleasant words and outgroup exemplars are paired with unpleasant words (*congruent trials*) rather than the reverse, then they display an implicit preference for their ingroup relative to the outgroup. If, however, individuals respond faster when ingroup exemplars are paired with unpleasant words and outgroup exemplars are paired with pleasant words (*incongruent trials*) rather than the reverse, then they display an implicit preference for the outgroup relative to the ingroup.

Although the IAT has fared well in terms of known groups validity (e.g., Cunningham, Preacher, & Banaji, 2001), less work has been conducted to examine whether variability in IAT scores are predictive of other variables. Results of studies that have been conducted on this issue

1. In addition to the studies we summarize in the text, Fazio et al. (1995) included black as well as white participants in their measure of implicit racial associations. However, there were only 8 black participants in this research. Additionally, Judd, Park, Ryan, Brauer, and Kraus (1995) assessed blacks' implicit biases, but their focus was on stereotype content rather than on relative preference for whites versus blacks. These studies, as well as one study of Livingston (2002, Study 1), used priming procedures to assess implicit associations. The remaining studies with black participants have used the Implicit Association Test (IAT).

are mixed (see Fazio & Olson, 2003), with some supporting the predictive utility of the IAT (e.g., McConnell & Liebold, 2001) and others not doing so (e.g., Karpinski & Hilton, 2001). Consequently, the IAT has uncertain validity at this time and, in a related vein, its interpretation has been questioned (Brendl, Markman, & Messner, 2001; Karpinski & Hilton, 2001). The present research thus contributes in a broad sense to the growing literature that seeks to understand the predictive utility of the IAT, as well as to our more specific interest in determining whether blacks' IAT performance is related to explicit attitudes and expressed preference for a white versus black interaction partner.

THE PRESENT RESEARCH: EXTENDING OUR UNDERSTANDING OF BLACKS' IMPLICIT ASSOCIATIONS

To address the question of whether blacks' implicit biases have meaningful implications for their actual judgments, we examined the relation between black participants' IAT scores and their preference for a white versus black work partner. If the patterns observed among black participants on the IAT reflect system justification processes, then they should be related to preferences in domains in which whites stereotypically outperform blacks, such as intelligence. We constructed an intellectually engaging task and gave black participants the opportunity to rate two potential partners, one of which was black and the other of which was white. In addition, we manipulated participants' "stakes" in the task by informing half of them that they and their partner stood to win \$100 each if they had the highest score on the joint task upon completion of the entire study. Although the idea of stakes has not been empirically tested, Jost and his colleagues suggest that system justification is most likely to occur when outcomes are important (Jost & Banaji, 1994). Hence, any relation between the IAT and partner preference might be particularly strong when participants are highly motivated to perform well.

We also collected various explicit measures to determine whether explicit attitudes play into partner preference and are related to IAT scores. To assess attitudes toward the ingroup, we borrowed from the Multidimensional Inventory of Black Identity (MIBI; Sellers, Smith, Shelton, Rowley, & Chavous, 1998). The MIBI makes a distinction between blacks' own personal views of their race (i.e., *private regard*) and their beliefs about how others view their race (i.e., *public regard*). It also captures the degree to which race is included in blacks' self-concepts (i.e., *central-*

ity). The MIBI includes other subscales that assess blacks' ideology regarding the way blacks should behave in society. These subscales, while interesting, have received less empirical attention than private and public regard and centrality and were not of theoretical interest in the present research. Given previous findings, we expected scores on private regard and centrality to be high and to vary little (e.g., Rowley et al., 1998). In contrast, we expected blacks' public regard scores to show greater variability, as blacks tend to differ in the extent to which they believe whites perceive them negatively (e.g., Johnson & Lecci, 2003; Livingston, 2002).

In a similar vein, we assessed blacks' attitudes toward whites with the recently developed Johnson–Lecci Scale (JLS; Johnson & Lecci, 2003). This measure is, to our knowledge, the only current measure of blacks' attitudes toward whites that draws upon blacks' unique experiences and was not developed by reversing race in measures designed primarily for whites. As Monteith and Spicer (2000) demonstrated, blacks' racial attitudes do not necessarily have the same basis as whites' racial attitudes. Although the JLS is relatively new, its authors report five studies that firmly establish its reliability and validity. The scale consistently yields four factors that concern discriminatory expectations of whites, negative beliefs about whites, negative views of intergroup relations, and negative verbal expressions toward whites.

Finally, we included an ideological measure that has been shown to be related to ingroup attitudes: social dominance orientation (SDO; e.g., Pratto, Sidanius, Stallworth, & Malle, 1994). SDO concerns the extent to which one favors group-based hierarchies, stratification, and domination (e.g., Levin & Sidanius, 1999; Sidanius, 1993). A recent investigation of the SDO scale (Jost & Thompson, 2000) revealed the existence of two distinct factors, opposition to equality (OEQ) and group-based dominance (GBD). GBD reflects one's views about the appropriateness of some groups dominating or being superior to other groups (e.g., "Superior groups should dominate inferior groups"). OEQ is focused more on whether equality is valued and desirable (e.g., "We should strive to make incomes more equal"). Jost and Thompson argued that for high-status groups OEQ and GBD are "two sides of the same coin" (p. 213). Consistent with this idea, the two factors were highly correlated for their white participants. In contrast, for low-status groups, GBD becomes ambiguous because agreement with group-based dominance can mean support for the ingroup or support

for other groups having dominance over the ingroup. In other words, unlike for whites, GBD is not simply the “flip side” of OEQ. Consistent with this reasoning, the correlation between GBD and OEQ items among black participants in the Jost and Thompson research was much lower than it was among white participants. Among black participants, therefore, we expected OEQ ratings to be a better indicator of system justifying tendencies and, in turn, to be more highly related to implicit bias than GBD ratings.

METHOD

PARTICIPANTS

Eighty-three (43 female and 40 male) black American undergraduates at the University of Kentucky volunteered to participate in the study. All participants were recruited by phone. Potential participants' names and phone numbers were obtained from the introductory psychology participant pool mass screening session and the registrar's list of black student enrollment. Twenty-eight participants were recruited from Introductory Psychology and compensated with research credit for their course grade. Fifty-five participants were recruited from the registrar's list and compensated with \$15.

IAT data for one participant were missing due to computer error. Partner preference data for three participants were missing because participants did not complete the scale as instructed (i.e., one participant drew a line between the “4” and “5” on each item as if to indicate no preference, one wrote in “4.5” for each item, and one provided a response on only one item and stated that the rest were irrelevant). Questionnaire data were missing for four participants: three are due to experimenter error, and one participant arrived late and did not have time to complete the questionnaires. Because of these missing data, degrees of freedom vary slightly across analyses.

APPARATUS

The IAT was programmed and administered on a personal computer using FIAT for Windows 2.3 (Farnham, 1998). Participants sat approximately 61 cm from the monitor while completing the IAT.

PROCEDURE

Participants completed the study individually. After obtaining informed consent, a white female experimenter explained (as part of the cover story) to participants that they would be participating in a study conducted by industrial organizational psychologists interested in assessing individual versus team performance. To enhance the cover story, participants were told that there were two other individuals participating simultaneously in nearby labs and that one participant would complete the tasks individually while the other two would be paired. Participants then completed a demographic questionnaire (e.g., name, sex, race, age, GPA) and believed the other two supposed participants did the same.

A rigged drawing ensured that all participants were assigned to the "team" condition and would work with a partner rather than individually. The experimenter explained that participants would be allowed to choose their partner based on the demographic information of the supposed other two participants. While the experimenter was allegedly interacting with these supposed others, participants were given time to practice a task that the experimenter described as similar to the upcoming team task. Specifically, participants were given 10 difficult anagrams to solve, which served to highlight the intellectual rigor of the task.

Upon the experimenter's return, participants were given time to read the demographic sheets of the supposed other participants. Each participant received a completed questionnaire from a black (James or Janet) and a white (Chris or Christy) person who were the same sex as the participant. The completed questionnaires included demographic information about the alleged other participants, including information such as age (18 or 19), year in school (both were described as freshmen), major (sociology or history), GPA (3.0 or 3.1), and hometown (Covington, KY or Bowling Green, KY). Responses were equated so that the primary difference between the two potential partners was their race.

At this point participants received one of two sets of instructions, determined randomly, as a manipulation of stakes. Specifically, half were told to choose a partner with whom they would feel comfortable working for the remaining time (low stakes condition). The other half were given this information, but were additionally told that each member of the team whose scores on the joint task were highest upon completion of the study would be awarded \$100 (high stakes condition). In actuality,

all participants were entered into a lottery and one randomly selected person was awarded the cash prize at the conclusion of data collection for the entire experiment.

Participants then rated their potential partners on the following five 8–point questions regarding *partner preference*: “To what extent would you get along better with one partner over the other?”; “To what extent are you more opposed to working with one partner over the other?”; “To what extent is it more important for you to work with one partner over the other?”; “To what extent do you think one partner is friendlier than the other?”; “To what extent do you think one partner is better at doing anagrams?” The alleged other participants’ names served as the anchors for each item (e.g., 1 = get along better with Partner A, Christopher; 8 = get along better with Partner B, James). In addition, whether a black versus white name appeared as the left versus right anchor was randomized. These items assessed participants’ overall positivity toward one partner over the other.² As a check for the stakes manipulation, participants rated the extent to which they were motivated to choose a successful partner (1 = not at all motivated to 7 = extremely motivated).

IAT. After completing this measure, participants were asked to help pilot a “new” measure, the IAT, while the experimenter allegedly set up the room in which the joint task would take place. Participants completed one of two programs that were created in order to control for trial order effects. These programs were identical except for the order in which participants completed the congruent (i.e., black name + pleasant word / white name + unpleasant word) and incongruent (i.e., black name + unpleasant word / white name + pleasant word) trial blocks. Participants were instructed to respond as quickly as possible while making as few mistakes as possible. The experimenter left the room while participants completed this task.

The IAT computer program began with general instructions that explained that participants would be assigning words to categories. The program then presented the categorization task, which contained seven blocks of trials. In each block, the categories (i.e., black, white, pleasant,

2. Participants were also asked to write the name of the partner with whom they would like to work, providing a dichotomous measure of partner choice. Although the partner preference ratings and partner choice measure were significantly and strongly correlated, $r(75) = .58, p < .01$, no other significant results were found for partner choice. The dichotomous partner choice item likely was not a sensitive enough measure.

unpleasant) were presented before the onset of the trials, and they remained on the left and right sides of the screen throughout the block. The stimulus words, which were obtained from Greenwald et al.'s (1998) stimulus lists, appeared in the center of the screen, one at a time in random order. Participants were instructed to assign each stimulus word as quickly as possible to the left or right category by pressing either the "A" key on the left side of the keyboard or the "5" key on the number pad, respectively. Correct categorizations were followed by green circles that appeared below the stimulus word, whereas incorrect responses were followed by red X's that remained on the screen until participants made the correct response. The inter-trial stimulus interval was 150 ms.

The first three blocks of trials were for practice. In the first block, participants categorized names as black versus white. In the second block, participants categorized words as pleasant versus unpleasant. In the third block, the tasks were combined; that is, participants were presented with both names and words and had to assign them to their appropriate categories. Two of the categories (e.g., black name + pleasant word) appeared together on the left and two (e.g., white name + unpleasant word) appeared on the right. This pairing was retained in the fourth block, but those trials served as test—rather than practice—trials. In the fifth (practice) block, the category placement was switched from that of the first block so that the category that had appeared on the left in the first block appeared on the right in the fifth block. This switch was maintained throughout the sixth and seventh blocks. The sixth (practice) block combined the categories (e.g., black name + unpleasant word on the left; white name + pleasant word on the right). The final block was identical to the sixth, but counted as a test block. Test blocks consisted of 40 trials each, and practice blocks consisted of twenty trials each.

When participants finished the IAT, the experimenter returned and informed participants that this part of the study was over and that they would not be working on a joint task with a partner. At this time, participants were asked whether they would be willing to participate in a second, supposedly unrelated study on ethnic attitudes.

Explicit Measures. After they provided informed consent, participants were given a packet containing the following questionnaires assembled in a random order. To assess racial identity, participants completed the MIBI (Sellers et al., 1998), a 56-item inventory scored on a 7-point Likert-type scale (1 = strongly disagree to 7 = strongly agree). The MIBI

taps various aspects of racial identity, but those most important for the purposes of the present research were private regard (e.g., “I feel good about black people”), public regard (e.g., “Overall, blacks are considered good by others”), and centrality (e.g., “In general, being black is an important part of my self-image”).

Attitudes toward whites were assessed using the 20-item JLS (Johnson & Lecci, 2003). Participants made responses on a 7-point Likert-type scale (1 = strongly disagree to 7 = strongly agree). The JLS includes items that tap blacks’ attitudes toward whites (e.g., “I consider myself to be racist toward whites”), perceptions of whites’ beliefs about blacks (e.g., “I believe that most whites really believe that blacks are genetically inferior”), and past behaviors regarding whites (e.g., “I have insulted a white person”).

SDO was assessed using the 16-item SDO scale (Pratto et al., 1994). Participants responded to the items using 7-point Likert-type scales to indicate their feelings toward each item (1 = very negative to 7 = very positive). The SDO scale included both OEQ (e.g., “We should do what we can to equalize conditions for different groups”) and GBD items (e.g., “To get ahead in life, it is sometimes necessary to step on other groups”).

Participants were instructed to place their completed packets in an envelope to help ensure confidentiality. Upon completion of the measures, the experimenter returned and probed participants for suspicion. We particularly wanted to determine whether participants saw a connection between the IAT and the partner-rating task and between the first part of the study and the questionnaires. No participant verbalized a link between these features of the study in any way. Finally, the experimenter debriefed, thanked, and compensated participants.

RESULTS

IAT EFFECT

FIAT for Windows 2.3 (Farnham, 1998) automatically drops the first two trials in each block. These initial latencies tend to be longer as participants are growing accustomed to the task. Similarly, FIAT recodes all response latencies that are less than 300 ms or greater than 3000 ms as 300 ms and 3000 ms, respectively. These procedures help ensure the validity of the task by eliminating extremely short and long response times that may be due to participants’ momentary inattention. The response laten-

cies were also log-transformed, given that reaction time data tend to be skewed (e.g., Ratcliff, 1993). Analyses were performed using the transformed data, but results are presented in milliseconds for ease of interpretation.

The response latencies were then analyzed in a 2 (participant sex: male vs. female) \times 2 (IAT order: congruent first vs. incongruent first) \times 2 (IAT trial type: congruent vs. incongruent) mixed-model analysis of variance (ANOVA), with repeated measures on the last factor. Results indicated a significant effect of IAT trial type only, such that participants took longer on the congruent ($M = 889.04$ ms) than incongruent trials ($M = 811.17$ ms), $F(1, 78) = 18.46, p < .001, d = .49$. Thus, participants overall displayed a moderate degree of implicit outgroup favoritism (i.e., more easily associating pleasant concepts with white names and unpleasant with black names).

IAT scores were then calculated by subtracting mean congruent trial latencies from incongruent latencies (see, for example, Greenwald et al., 1998), such that positive scores reflect ingroup favoritism and negative scores reflect outgroup favoritism. Descriptive statistics for the IAT and all other measures are presented in Table 1, and the distribution of IAT scores is depicted in Figure 1. Approximately 40% of the sample displayed implicit ingroup favoritism, responding significantly faster when black names were paired with pleasant words and white names were paired with unpleasant words than the reverse. More strikingly, 60% of the sample displayed implicit outgroup favoritism, responding significantly faster when white names were paired with pleasant words and black names with unpleasant words than the reverse. Indeed, overall the IAT effect was significantly different from zero and in a negative direction, $t(81) = 3.81, p < .001$, underscoring the degree to which many black participants in our sample exhibited relatively negative ingroup associations. On one hand, the present findings replicate those of previous studies (e.g., Livingston, 2002) in the high degree of variability among blacks on the IAT, with some participants exhibiting implicit ingroup favoritism and some exhibiting implicit outgroup favoritism. On the other hand, the significant degree of outgroup favoritism in the present study is a departure from previous findings. Specifically, Livingston (2002) found no evidence of ingroup or outgroup preference among black participants in two samples. We will return to this point in the General Discussion.

TABLE 1. Descriptive Statistics for All Measures

	α	Minimum	Maximum	<i>M</i>	<i>SD</i>
IAT	—	-561.00	641.00	-77.87	211.85
Private regard	.70	5.17	7.00	6.73	.42
Public regard	.63	1.67	5.17	3.61	.89
Centrality	.71	2.75	6.88	5.25	1.00
JLS	.92	1.40	5.75	3.38	1.18
GBD	.76	1.00	4.38	1.87	.88
OEQ	.77	1.00	4.00	1.59	.65
SDO	.77	1.00	4.06	1.73	.60
Partner preference	.69	2.50	7.25	4.77	.85

Note. Descriptives for IAT are provided in milliseconds. All other measures, except partner preference, are on 7-point scales. Partner preference is on an 8-point scale, with higher numbers indicating a greater preference for the black partner.

EXPLICIT MEASURES

All scales were computed by reverse-scoring items when appropriate and averaging across all items to create indexes ranging from 1 – 7, with higher scores indicating greater levels of endorsement for each construct. Descriptive statistics for each scale are presented in Table 1. What is noteworthy about the explicit measures is that they are less variable and reflect more positive feelings toward the ingroup than participants' IAT scores, consistent with previous research (e.g., Livingston, 2002).

PARTNER PREFERENCE

To construct an index of partner preference, the five 8-point items were reverse-scored when appropriate such that higher numbers indicated a greater preference for a black partner. The items were first submitted to a reliability analysis. With all five items included, the reliability coefficient was rather low ($\alpha = .58$). The analysis indicated, however, that the reliability could be substantially improved with the removal of one item ("To what extent are you more opposed to working with one partner over the other?"). Further, whereas the other four items correlated well with each other (average inter-item $r = .37$), they correlated rather poorly with this item (average inter-item $r = -.03$). To verify that this item stood alone, the items were then submitted to a principal compo-

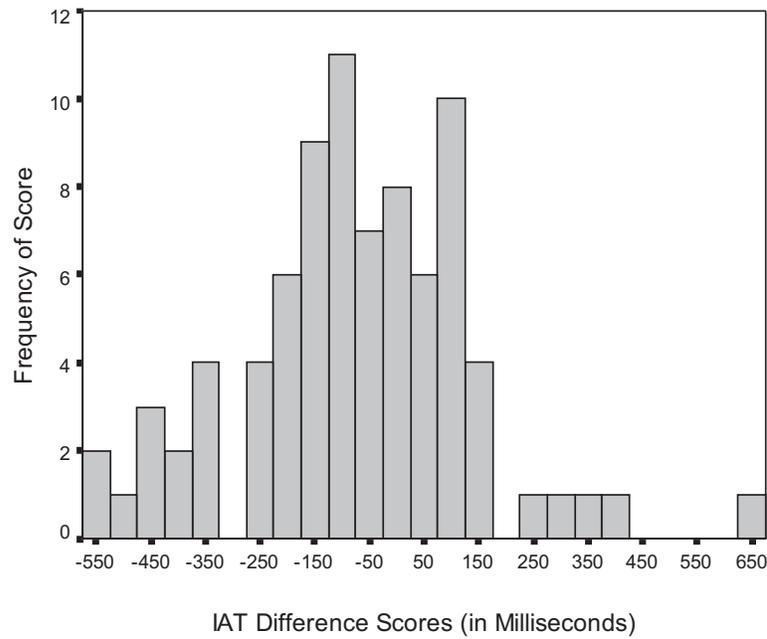


FIGURE 1. Distribution of IAT scores. Negative IAT scores indicate implicit preference for whites relative to blacks. Positive scores indicate implicit preference for blacks relative to whites.

ment analysis with varimax rotation. Two factors emerged that accounted for approximately 66% of the variance. Indeed, the one item appeared distinct from the others, and it loaded on its own factor. Hence, this item was excluded, yielding a 4-item index of partner preference with substantially improved reliability, as shown in Table 1.

PREDICTING PARTNER PREFERENCE

As illustrated in Figure 2, participants' IAT scores and their partner preferences were significantly and positively related, $r(77) = .23, p = .04$. The more blacks implicitly preferred their own race relative to whites, the greater their preference for a black relative to a white work partner. Importantly, this relationship suggests that blacks' implicit biases are more than latent associations; they do manifest in actual intergroup judgments.

TABLE 2. Zero-Order Correlations among All Measures

	1	2	3	4	5	6	7	8
1. IAT	—							
2. Private regard	.40**	—						
3. Public regard	.12	-.18	—					
4. Centrality	.23*	.55**	-.24*	—				
5. JLS	.13	.28*	-.39**	.46**	—			
6. GBD	.10	-.15	.27*	.04	.15	—		
7. OEQ	-.38**	-.37**	.13	-.17	-.19	.20	—	
8. SDO	-.13	-.31**	.27*	-.06	.01	.85**	.69**	—
9. Partner preference	.23*	.14	.13	.13	.32**	.28*	-.10	.15

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

Partner preference was also significantly correlated with some of the explicit measures, as shown in Table 2. First, partner preference was related to the JLS. The stronger blacks' anti-white attitudes were, the greater their preference for a black versus a white partner. In addition, partner preference and GBD were significantly correlated, such that the more blacks endorsed the idea of group dominance or superiority, the more they preferred a black versus white partner. It appears, then, that for black participants in this sample, GBD reflected a desire for ingroup, rather than outgroup, dominance (cf. Jost & Thompson, 2000).

Would the IAT remain a significant correlate of partner preference when controlling for the explicit measures with which preference was also related? To address this issue, a partial correlation coefficient was computed to examine the relationship between partner preference and IAT scores, controlling for JLS and GBD. Although the relationship was marginal, $r(71) = .21, p = .08$, its magnitude was maintained whether or not the explicit measures were partialled out. This finding suggests that, despite the somewhat larger correlations between the explicit measures and partner preferences, implicit biases nonetheless played an important, independent role in predicting actual judgments of blacks versus whites.

RELATIONSHIPS BETWEEN IAT AND EXPLICIT MEASURES

Although our primary focus in the research was to determine whether implicit biases were related to partner preference, we also wished to de-

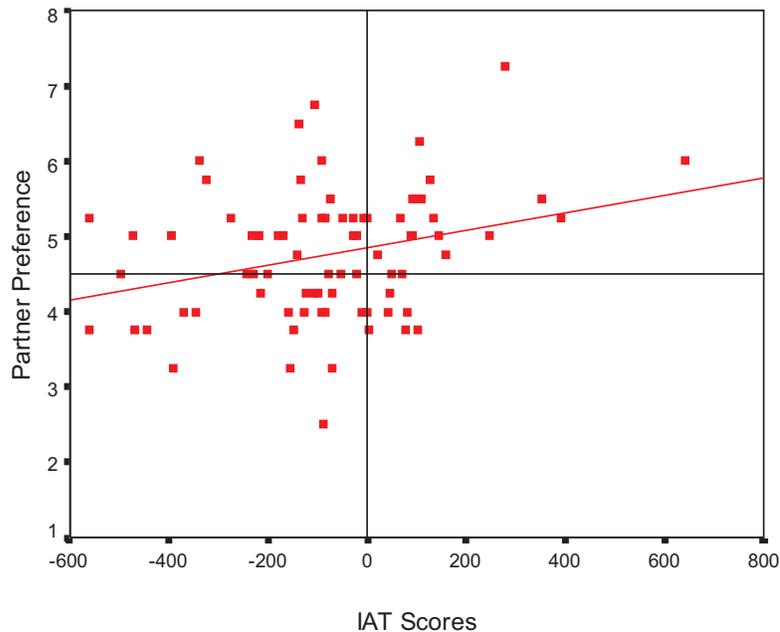


FIGURE 2. The relationship between IAT scores and partner preference. Partner preference scores above the midpoint indicate a preference for the black partner, whereas those below the midpoint indicate a preference for the white partner.

termine whether explicit attitudes were related to IAT scores. The zero-order correlations for all measures are reported in Table 2. The ranges of the explicit measures were relatively restricted (compared with the range of IAT scores), but significant correlations with the IAT nonetheless emerged. First, participants' implicit biases were significantly related to aspects of racial identity. The more positively they viewed their ingroup at the implicit level, the higher their explicit private regard for their race and the more central race was to their self-concept. There was, however, no relationship between participants' attitudes toward whites and their implicit biases. Taken together, these correlational findings suggest that blacks' implicit racial biases are not entirely dissociated from their consciously held racial attitudes. In addition, outgroup favoritism among blacks appears to have less to do with how they feel about whites than about their racial ingroup.

TABLE 3. Summary of Hierarchical Regression Analysis for Variables Predicting Partner Preference

Variable	<i>B</i>	<i>SE B</i>	β
Step 1			
Stakes	-0.19	0.19	-0.11
IAT	0.9	0.4	.25*
Step 2			
Stakes	-0.14	0.25	-0.08
IAT	0.77	0.59	0.22
Stakes \times IAT	0.24	0.81	0.05

Note. $N = 78$. * $p < .03$.

Surprisingly, the relationship between the IAT and public regard was not significant. This stands in contrast to findings obtained by Livingston (2002), in which blacks' implicit biases appeared to stem from their perceptions of whites' negativity toward their race. One possible explanation for this discrepancy is that the measure of public regard used in our study differed slightly from those used in Livingston's studies. In addition, given that the IAT is a relative measure, perhaps a better explicit indicator of public regard would be to ask blacks to report their perceptions of how others view blacks relative to whites.

Finally, blacks' implicit biases were related to system justifying ideology. Specifically, the more participants' scores reflected an outgroup favoritism bias, the stronger their opposition to equality (i.e., OEQ). This means that the more participants implicitly favored whites, the more they supported a hierarchical structure within society. Further, as previous research by Jost and Thompson (2000) would suggest, blacks' implicit biases were not related to GBD.

DO SITUATIONAL STAKES MODERATE THE IAT–PARTNER PREFERENCE RELATIONSHIP?

Recall that we suspected that the relationship between blacks' implicit biases and their actual judgments would be stronger when participants are especially invested in the outcome of their judgment (i.e., high stakes condition). To ensure that the manipulation of stakes was successful, participants' scores on the single item regarding their motivation to

choose a successful partner were submitted to an independent groups *t*-test. Indeed, participants in the low stakes condition reported significantly less motivation ($M = 3.95$) than those in the high stakes condition ($M = 4.93$), $t(81) = 2.68$, $p < .01$, $d = .60$. It is important to note, however, that even participants in the low stakes condition were quite motivated to choose a successful partner, scoring near the midpoint of the scale.

Partner preference was regressed on stakes condition (dummy-coded as 0 = low stakes, 1 = high stakes), IAT scores (which were centered, following Aiken & West, 1991) and their interaction. Main effects were entered on the first step, followed by the two-way interaction on the second. All effects were interpreted on the step at which they were entered. As shown in Table 3, the only significant finding to emerge was a main effect of IAT scores. The greater participants' outgroup favoritism tendencies the stronger their preference for a white versus a black partner. Stakes had no effect, either direct or moderating, on partner preference. Hence, the relationship between implicit bias and partner preference generalized across conditions. We suspect that this is because even participants in the low stakes condition were, as previously noted, quite motivated to work with the partner they believed was more likely to succeed at the task.

We also conducted a regression analysis substituting participants' centered scores on the manipulation check item for their randomly assigned condition. Results of this analysis indicated a direct, but not moderating, effect of participants' motivation to choose a successful partner on their partner preference, $F(1, 76) = 4.59$, $p < .04$, $\beta = -.24$. That is, the more motivated participants were to choose a successful partner, the greater their preference for a white work partner. Importantly, participants' implicit biases still independently predicted their partner preference even when controlling for their motivation, $F(1, 76) = 6.04$, $p < .02$, $\beta = .27$.³

3. Partner preference was also regressed on condition, explicit measure (centered, following Aiken & West, 1991), and their interaction in separate analyses for each of the explicit measures. Main effects were entered on the first step, followed by the two-way interaction on the second. All effects were interpreted on the step at which they were entered. Results paralleled the correlations provided in Table 2; that is, only JLS and GBD were significant predictors of partner preference. None of the analyses yielded significant effects of stakes condition or its interaction with the explicit measures.

DISCUSSION

As Allport (1954) suggested a half-century ago, black Americans have so long been exposed to stereotypes “that they are lazy, ignorant, dirty, and superstitious that they may half believe the accusations” (p. 152). His statement may have been a foreshadowing of the patterns of bias revealed by implicit measures. Indeed, mounting evidence suggests that many—though not all—black individuals do hold outgroup-favoring associations at the implicit level (Livingston, 2002; Nosek et al., 2002; Spicer, 1999). The present research demonstrates that the associations not only appear to be widely held, but they also have important implications.

CONTRIBUTIONS OF THE PRESENT RESEARCH

First, we demonstrated a relationship between blacks’ implicit biases and actual intergroup judgments in a context that blacks could easily encounter in the everyday world. Previous research has shown that implicit outgroup favoritism (or at least an attenuated ingroup bias) exists (Livingston, 2002; Nosek et al., 2002; Spicer, 1999), and some research has offered compelling reasons for why such a bias exists (Livingston, 2002). The present research adds to the extant literature by showing that these associations are not mere cognitive structures that reflect societal patterns or hierarchies, but that the associations can translate into meaningful intergroup judgments. This finding potentially has implications for real world situations in which race may factor into important decisions.

In addition, with the present research we are beginning to unveil some of the factors associated with individual differences in blacks’ implicit racial biases. This is no minor point, because so little literature exists on this topic and because we know from the literature on explicit attitudes that blacks’ racial attitudes may not have the same basis as whites’ (Monteith & Spicer, 2000). System justification theory (Jost & Banaji, 1994) suggests that such implicit biases may stem from a motivation to legitimize the hierarchy observed within society. Indeed, we found that blacks’ implicit biases were related to ideology consistent with system justification. More germane to the present research, work by Livingston (2002) demonstrated that blacks’ implicit biases are related to perceived negativity from whites; thus, he argued that these biases reflect an internalization of racism, in the tradition of Allport (1954). We also believe that this is likely the case, but the present findings suggest that such bi-

ases may stem from blacks' relatively less positive beliefs about their ingroup rather than from some liking or admiration for whites (i.e., the IAT was related to private regard and centrality but not JLS). In this sense, we feel that our research both complements and extends that of Livingston (2002).

It appears that, although blacks' explicit racial attitudes have a different basis than whites' (Monteith & Spicer, 2000), the basis of implicit racial associations among many blacks may not differ from that of whites. Specifically, blacks' explicit racial attitudes are largely related to perceived racism (Johnson & Lecci, 2003; Monteith & Spicer, 2000) whereas whites' explicit racial attitudes are related to egalitarian ideals (Monteith & Spicer, 2000). In contrast, the basis of both blacks' and whites' implicit racial associations—when they are relatively negative toward blacks—may be exposure to stereotypic images (see Devine, 1989) and to the ever-present reminders of the social hierarchy in which whites are held in higher esteem (Jost & Banaji, 1994).

The present findings also add to the growing body of literature regarding the validity of implicit biases as measured by the IAT. Despite mixed reviews (e.g., Brendl et al., 2001; Karpinski & Hilton, 2001), the IAT has exhibited relationships with various attitudes, behaviors, and judgments that one would expect. For example, in research with white participants, there is often a moderate-sized correlation between IAT scores and explicit attitudes toward blacks (e.g., Cunningham et al., 2001; McConnell & Liebold, 2001; Monteith, Voils, & Ashburn-Nardo, 2001); that is, the greater white participants' consciously held anti-black attitudes, the more racially biased their implicit associations. In addition, recent investigations have demonstrated that the IAT does have the ability to predict behavior successfully. For example, McConnell and Liebold (2001) found that white participants' IAT scores predicted certain avoidance behaviors toward a black versus a white experimenter. In addition, Nosek, Banaji, and Greenwald (2002) demonstrated that females' implicit attitudes toward math were correlated with their performance on the quantitative portion of the Scholastic Aptitude Test. Interestingly, IAT scores sometimes even predict behaviors that explicit measures do not (McConnell & Liebold, 2001; Rudman & Glick, 2001).

In the present research, IAT scores were predictive of judgments in a context that one would expect: a domain in which whites

stereotypically are expected to excel. Although our suspicion was that implicit biases would be especially likely to surface when stakes were high, it appears that our manipulation of stakes fell slightly short of our expectations; that is, even participants in the low stakes condition were quite motivated to choose a successful partner. Nonetheless, we still believe that implicit biases are more likely to manifest in judgments or behaviors in some situations than in others. Blacks with strong outgroup–favoring biases may, for example, be less likely to put their trust in a black versus a white surgeon, or to value the wisdom of a black versus a white professor, or to endorse a black versus a white Presidential candidate. The common thread in these situations is that there are strong stereotypes about blacks' versus whites' ability or competence. In a social situation, however, we suspect that blacks' implicit biases would have little bearing on their preferences for a white versus a black companion.

LIMITATIONS AND ALTERNATIVE EXPLANATIONS

In the present study, we examined the relationships between the IAT and measures that were theoretically plausible. This is not to say that other constructs, unexamined in the present research, would not better account for the variability in blacks' or others' implicit biases. With future research that more thoroughly explicates the factors that account for the variability in blacks' implicit biases more specific predictions regarding the conditions under which such biases are likely to predict judgments and behaviors can be made. In an ongoing study in our lab, we are examining the relationship between sociocultural factors (e.g., childhood socialization and intergroup contact) and blacks' IAT scores to provide a more complete picture. In addition, in the present study we chose a single domain in which status differences would be made salient, and we did so based on system justification theory, one of the few theories that speaks directly to the effects that social hierarchies (in which, for example, one race is held in higher esteem than another) have on members of lower status groups. There are likely other domains in which such biases would manifest. However, given the relationship between IAT scores and system justifying ideology in the present research, we are confident that such motivational processes were involved in this context to some degree.

There is, however, mixed support for system justification theory in the extant literature—particularly in research employing black participants. Whereas the present research yielded a significant degree of outgroup favoritism, previous studies have instead yielded more neutral patterns that indicate neither ingroup nor outgroup favoritism overall (Livingston, 2002). Further research is needed to determine why system-justifying processes seem to operate more in some social groups than in others and why some group members are more vulnerable than others.

One might question our use of the IAT, given its various criticisms. As noted previously, we believed that the IAT was well suited for our purposes because it provides a relative index of racial preference. Drawing from the system justification literature, we were most interested in situations in which blacks might express a preference for whites relative to blacks, and we felt that the IAT captures this distinction somewhat better than other implicit measures (e.g., priming). Of course, the IAT does not allow for the assessment of purely positive or negative associations with either the ingroup or outgroup; indeed, its dual categorization nature does not allow one to separate ingroup/outgroup favoritism from derogation. Other implicit measures, such as priming procedures (e.g., Fazio et al., 1995; Gaertner & McLaughlin, 1983; Payne, 2001; Perdue, Dovidio, Gurtman, & Tyler, 1990; Wittenbrink et al., 1997) and the go/no-go association test (Nosek & Banaji, 2001) offer that advantage. Note, however, that some research suggests that the IAT performs similarly to other implicit measures (e.g., Cunningham et al., 2001), even among black participants (Livingston, 2002). In future research, perhaps multiple implicit measures would not only provide the most thorough assessment of implicit biases, but they would also supplement the developing literature on the validity of the IAT.

Future research will also need to assess actual behavior, rather than focusing on reported preferences. The measure that we used that came closest to assessing actual behavior was the dichotomous partner choice measure, and partner choice was not associated with any significant results (see Footnote 2). This may be because the dichotomous partner choice was not sensitive enough; alternatively, perhaps blacks' implicit racial biases are unrelated to actual behavior.

Most importantly, we recognize that the present findings cannot provide a firm basis for causality. We did not counterbalance the order of

the experimental tasks for fear of rousing suspicion. We chose this specific task order because we believed that it was less likely for partner preference to influence IAT scores than for IAT scores to influence partner preference. Although implicit biases can be manipulated in various contexts (e.g., Blair, Ma, & Lenton, 2001; Dasgupta & Greenwald, 2001; Kawakami, Dovidio, Moll, Hermsen, & Russin, 2000), this typically occurs after repeated exposure to multiple counterstereotypic images or repeated instructions to negate stereotypes or visualize counterstereotypic exemplars. Thus, we suspect that it would take more than a rating of partner preference to influence IAT scores. Nevertheless, self-perception theory (Bem, 1972) would suggest that participants' behavior (i.e., partner preference) would influence their attitudes, as participants would be motivated to rationalize their recent preference. Though possible, we suspect that this possibility is not very probable. Specifically, IAT scores were related not only to partner preference but also to the theoretically relevant explicit measures of private regard, racial centrality, and opposition to equality. Because these explicit measures were not related to partner preferences, we think it is unlikely that IAT performance and performance on the explicit measures were being determined by recent judgments of partner preference. However, future research should consider either separate sessions for testing attitudes and behavior or counterbalancing the measures.

CONCLUSION

Our results indicated that to the extent that blacks exhibited an implicit evaluative preference for whites over blacks, the greater their acceptance of social inequalities and the weaker their sense of confidence in a black than a white partner's chances of success on an intellectually challenging task. Although partner preferences overall were not extreme, the fact that they were related to biases held by so many may unfortunately serve to perpetuate inequality in situations in which members of traditionally disadvantaged groups may be able to gain positions of power or leadership. For this reason, we believe it is critical that researchers continue to explore the consequences of variation in implicit racial biases among blacks. And, perhaps even more importantly, research is needed to understand what leads to this variation.

REFERENCES

- Aiken, L. S., & West, S. G. (1991). *Multiple regression: Testing and interpreting interactions*. Newbury Park, CA: Sage Publications.
- Allport, G. W. (1954). *The nature of prejudice*. Reading, MA: Addison-Wesley.
- Ashburn-Nardo, L., Voils, C. I., & Monteith, M. J. (2001). Implicit associations as the seeds of intergroup bias: How easily do they take root? *Journal of Personality and Social Psychology, 81*, 789–799.
- Banse, R., Seise, J., & Zerbes, N. (2001). Implicit attitudes towards homosexuality: Reliability, validity, and controllability of the IAT. *Zeitschrift fuer Experimentelle Psychologie, 48*, 145–160.
- Bargh, J. A., Chen, M., & Burrows, L. (1996). Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action. *Journal of Personality and Social Psychology, 71*, 230–244.
- Bem, D. J. (1972). Self-perception theory. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 6, pp. 1–62). New York: Academic Press.
- Blair, I. V., Ma, J. E., & Lenton, A. P. (2001). Imagining stereotypes away: The moderation of implicit stereotypes through mental imagery. *Journal of Personality and Social Psychology, 81*, 828–841.
- Brendl, C. M., Markman, A. B., & Messner, C. (2001). How do indirect measures of evaluation work? Evaluating the inference of prejudice in the Implicit Association Test. *Journal of Personality and Social Psychology, 81*, 760–773.
- Crocker, J., Luhtanen, R., Blaine, B., & Broadnax, S. (1994). Collective self-esteem and psychological well-being among white, black, and Asian college students. *Personality and Social Psychology Bulletin, 20*, 503–513.
- Cunningham, W. A., Preacher, K. J., & Banaji, M. R. (2001). Implicit attitude measures: Consistency, stability, and convergent validity. *Psychological Science, 12*, 163–170.
- Dasgupta, N., & Greenwald, A. G. (2001). On the malleability of automatic attitudes: Combating automatic prejudice with images of admired and disliked individuals. *Journal of Personality and Social Psychology, 81*, 800–814.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56*, 5–18.
- Dovidio, J. F., Kawakami, K., Johnson, C., Johnson, B., & Howard, A. (1997). On the nature of prejudice: Automatic and controlled processes. *Journal of Experimental Social Psychology, 33*, 510–540.
- Egloff, B., & Schmukle, S. C. (2002). Predictive validity of an Implicit Association Test for assessing anxiety. *Journal of Personality and Social Psychology, 83*, 1441–1455.
- Farnham, S. D. (1998). The Farnham Implicit Association Test for Windows (Version 2.3) [Computer software]. University of Washington.
- Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology, 69*, 1013–1027.
- Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology, 54*, 297–327.
- 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.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology, 74*, 1464–1480.

- Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, 30, 513–541.
- Johnson, J. D., & Lecci, L. (2003). Assessing anti-white attitudes and predicting perceived racism: The Johnson–Lecci Scale. *Personality and Social Psychology Bulletin*, 29, 299–312.
- Jost, J. T., & Banaji, M. R. (1994). The role of stereotyping in system-justification and the production of false consciousness. *British Journal of Social Psychology*, 33, 1–27.
- Jost, J. T., Pelham, B., & Carvallo, M. (2002). Non-conscious forms of system justification: Implicit and behavioral preferences for higher status groups. *Journal of Experimental Social Psychology*, 38, 586–602.
- Jost, J. T., & Thompson, E. P. (2000). Group-based dominance and opposition to equality as independent predictors of self-esteem, ethnocentrism, and social policy attitudes among African Americans and European Americans. *Journal of Experimental Social Psychology*, 36, 209–232.
- Judd, C. M., Park, B., Ryan, C. S., 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.
- Karpinski, A., & Hilton, J. L. (2001). Attitudes and the Implicit Association Test. *Journal of Personality and Social Psychology*, 81, 774–788.
- Kawakami, K., Dovidio, J. F., Moll, J., Hermsen, S., & Russin, A. (2000). Just say no (to stereotyping): Effects of training in the negation of stereotypic associations on stereotype activation. *Journal of Personality and Social Psychology*, 78, 871–888.
- Kim, D.-Y., & Greenwald, A. G. (1998, May). Voluntary controllability of implicit cognition: Can implicit attitudes be faked? Paper presented at the annual meeting of the Midwestern Psychological Association, Chicago.
- Lepore, L., & Brown, R. (1997). Category and stereotype activation: Is prejudice inevitable? *Journal of Personality and Social Psychology*, 72, 275–287.
- Levin, S., & Sidanius, J. (1999). Social dominance and social identity in the United States and Israel: Ingroup favoritism or outgroup derogation? *Political Psychology*, 20, 99–126.
- Livingston, R. W. (2002). The role of perceived negativity in the moderation of African Americans' implicit and explicit racial attitudes. *Journal of Experimental Social Psychology*, 38, 405–413.
- McConnell, A. R., & Liebold, J. M. (2001). Relations among the Implicit Association Test, discriminatory behavior, and explicit measures of racial attitudes. *Journal of Experimental Social Psychology*, 37, 435–442.
- Monteith, M. J., & Spicer, C. V. (2000). Contents and correlates of whites' and blacks' racial attitudes. *Journal of Experimental Social Psychology*, 36, 125–154.
- Monteith, M. J., Voils, C. I., & Ashburn–Nardo, L. (2001). Taking a look underground: Detecting, interpreting, and reacting to implicit racial biases. *Social Cognition*, 19, 395–417.
- Mullen, B., Brown, R. J., & Smith, C. (1992). Ingroup bias as a function of salience, relevance, and status: An integration. *European Journal of Social Psychology*, 22, 103–122.
- Nosek, B. A., & Banaji, M. R. (2001). The Go/No-go Association Task. *Social Cognition*, 19, 625–666.
- Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002). Math = male, me = female, therefore math ≠ me. *Journal of Personality and Social Psychology*, 83, 44–59.
- Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002). Harvesting implicit group attitudes and beliefs from a demonstration website. *Group Dynamics*, 6, 101–115.
- Payne, B. K. (2001). Prejudice and perception: The role of automatic and controlled pro-

- cesses in misperceiving a weapon. *Journal of Personality and Social Psychology*, 81, 181–192.
- Perdue, C. W., Dovidio, J. F., Gurtman, M. B., & Tyler, R. B. (1990). Us and them: Social categorization and the process of intergroup bias. *Journal of Personality and Social Psychology*, 59, 475–486.
- Pratto, F., Sidanius, J., Stallworth, L. M., & Malle, B. F. (1994). Social dominance orientation: A personality variable predicting social and political attitudes. *Journal of Personality and Social Psychology*, 67, 741–763.
- Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin*, 114, 510–532.
- Rowley, S. J., Sellers, R. M., Chavous, T. M., & Smith, M. A. (1998). The relationship between racial identity and self-esteem in African American college and high school students. *Journal of Personality and Social Psychology*, 74, 715–724.
- Rudman, L. A., Feinberg, J., & Fairchild, K. (2002). Minority members' implicit attitudes: Automatic ingroup bias as a function of group status. *Social Cognition*, 20, 294–320.
- Rudman, L. A., & Glick, P. (2001). Prescriptive gender stereotypes and backlash toward agentic women. *Journal of Social Issues*, 57, 743–762.
- Schuman, H., Steeh, C., Bobo, L., & Krysan, M. (1997). *Racial attitudes in America: Trends and interpretations*. Cambridge, MA: Harvard University Press.
- Sekaquaptewa, D., Espinoza, P., Thompson, M., Vargas, P., & von Hippel, W. (2003). Stereotypic explanatory bias: Implicit stereotyping as a predictor of discrimination. *Journal of Experimental Social Psychology*, 39, 75–82.
- Sellers, R. M., Smith, M. A., Shelton, J. N., Rowley, S. A. J., & Chavous, T. M. (1998). Multidimensional model of racial identity: A reconceptualization of African American racial identity. *Personality and Social Psychology Review*, 2, 18–39.
- Sidanius, J. (1993). The psychology of group conflict and the dynamics of oppression: A social dominance perspective. In S. Iyengar & W. J. McGuire (Eds.), *Explorations in political psychology* (pp. 183–219). Durham, NC: Duke University Press.
- Spicer, C. V. (1999). *Effects of self-stereotyping and stereotype threat on African Americans' intellectual performance*. Unpublished doctoral dissertation, University of Kentucky, Lexington, Kentucky.
- Tajfel, H., Billig, M. G., Bundy, R. P., & Flament, C. (1971). Social categorization and intergroup behaviour. *European Journal of Social Psychology*, 1, 149–178.
- Vanman, E. J., Paul, B. Y., Ito, T. A., & Miller, N. (1997). The modern face of prejudice and structural features that moderate the effect of cooperation on affect. *Journal of Personality and Social Psychology*, 73, 941–955.
- von Hippel, W., Sekaquaptewa, D., & Vargas, P. (1997). The linguistic intergroup bias as an implicit indicator of prejudice. *Journal of Experimental Social Psychology*, 33, 490–509.
- Wittenbrink, B., Judd, C. M., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology*, 72, 262–274.