

Group Processes & Intergroup Relations

<http://gpi.sagepub.com/>

Discrimination and the Implicit Association Test

Laurie A. Rudman and Richard D. Ashmore

Group Processes Intergroup Relations 2007 10: 359

DOI: 10.1177/1368430207078696

The online version of this article can be found at:

<http://gpi.sagepub.com/content/10/3/359>

Published by:



<http://www.sagepublications.com>

Additional services and information for *Group Processes & Intergroup Relations* can be found at:

Email Alerts: <http://gpi.sagepub.com/cgi/alerts>

Subscriptions: <http://gpi.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations: <http://gpi.sagepub.com/content/10/3/359.refs.html>

Discrimination and the Implicit Association Test

Laurie A. Rudman

Rutgers University

Richard D. Ashmore

Rutgers University

Prejudice researchers have been criticized for failing to assess behaviors that reflect overtly hostile actions (i.e. racial animus; Arkes & Tetlock, 2004; Mackie & Smith, 1998). Two studies sought to begin to fill this gap in the implicit literature by showing that scores on the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998) are linked to harmful intergroup behaviors. In Study 1, the IAT predicted self-reported racial discrimination, including verbal slurs, exclusion, and physical harm. In Study 2, the IAT predicted recommended budget cuts for Jewish, Asian, and Black student organizations (i.e. economic discrimination). In each study, evaluative stereotype (but not attitude) IATs predicted behaviors even after controlling for explicit attitudes. In concert, the findings suggest that implicit stereotypes are more predictive of overtly harmful actions than implicit attitudes in the intergroup relations domain.

KEYWORDS discrimination, implicit prejudice, implicit stereotypes, racial stereotypes, social cognition, intergroup relations

EMPIRICAL observation distinguishes scientific psychologists from philosophers, novelists, and priests, who are also purveyors of truths about human nature. Because the quality of our insights depends on the clarity of our tools, methodological advances are the lifeblood of psychology. When a method is introduced, it should rightly be the target of skepticism and debate until its usefulness has been established. This was the expected trajectory for the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998), a response latency task that was developed to measure implicit attitudes, but has since proved useful for assessing other constructs, including implicit stereotypes (e.g. Greenwald, Pickrell, & Farnham, 2002).

The ingeniously simple concept underlying the IAT is that tasks are performed well when they rely on well-practiced associations between objects and attributes. In the attitude IAT, respondents categorize two classes of objects (e.g. dogs and cats) with both good and bad words (e.g. vacation vs. poison). An automatic preference for dogs is shown to the extent that

Author's note

Address correspondence to, Laurie A. Rudman, Department of Psychology, Tillett Hall, Rutgers, the State University of New Jersey, 53 Avenue E, Piscataway, New Jersey 08854-8040, USA
[email: rudman@rci.rutgers.edu]

the pro-dog task (dogs + good/cats + bad) is performed faster and more accurately than the pro-cat task (dogs + bad/cats + good). For stereotype IATs, good and bad words are replaced with specific attributes associated with each object (e.g. loyal vs. aloof). If the stereotype congruent task (dogs + loyal/cats + aloof) is performed faster and more accurately than the stereotype incongruent task (dogs + aloof/cats + loyal), an implicit stereotype is shown.

Although still young, the IAT has been tested in over 100 studies—far more so than any other response latency technique. A recent meta-analysis (Poehlman, Uhlmann, Greenwald, & Banaji, 2004) supported the IAT's temporal stability, internal consistency, and criterion validity (e.g. the IAT predicted voting, Scholastic Aptitude Test [SAT] scores, and consumer choice). Most promisingly, the IAT predicted behaviors better than self-reports did when the domain concerned prejudice and stereotypes. The behaviors included target evaluations, hiring decisions, and pro-social indicators (both verbal and nonverbal), suggesting a wide range of utility for the IAT. Given that self-reported prejudice was less useful when these behaviors were at stake, the IAT appears to be a promising methodological advance.

In addition, implicit associations behave in accord with classic attitude and intergroup theories (for a review, see Uhlmann & Poehlman, 2005). For example, they are sensitive to context and conditioning, just as attitudes and prejudice are (Blair, 2002; Fazio & Olson, 2003). Further, IAT scores have supported cognitive consistency principles (Greenwald, Pickrell, & Farnham, 2002), the contact hypothesis (Rudman, Ashmore, & Gary, 2001), aversive racism theory (Son Hing, Li, & Zanna, 2002), social identity theory (Ashburn-Nardo, Voils, & Monteith, 2001), and system justification theory (Jost, Pelham, & Carvallo, 2002; Rudman, Feinberg, & Fairchild, 2002). Finally, the relationship between implicit and explicit attitudes can be characterized as heterogeneous (Blair, 2001; Fazio & Olson, 2003; Nosek, Banaji, & Greenwald, 2002), but it is moderated by theoretically expected variables, including attitude strength, social desirability, and measurement error (Cunningham, Preacher,

& Banaji, 2001; Karpinski, Steinman, & Hilton, 2005; Nosek, 2005).

However, implicit researchers can be criticized—along with prejudice researchers in general—for rarely assessing overtly hostile behaviors (Mackie & Smith, 1998). This oversight affords an opening for implicit social cognition critics. In particular, the IAT's predictive utility in the prejudice domain has been questioned by Arkes and Tetlock (2004)—along with behavioral data for other implicit measures (e.g. Dovidio, Kawakami, & Gaertner, 2002; Fazio, Jackson, Dunton, and Williams, 1995)—on the basis that *some* behaviors might reflect emotions other than antipathy (e.g. nonverbal reactions, such as gaze or speech disfluencies, might indicate guilt or anxiety). However, in Poehlman et al.'s (2004) meta-analysis of IAT findings, less than a third of the behaviors assessed were nonverbal (or otherwise ambiguous). Yet, as already noted, the IAT was a better predictor of these behaviors, compared with self-reports. Moreover, the role of anxiety in prejudicial responding has long been recognized (e.g. Islam & Hewstone, 1993; Stephan & Stephan, 1985). That is, people may feel anxious in the presence of outgroup members, but this does not mean they are egalitarians. In fact, prejudice is typically defined as a negative orientation that can be expressed as moving against (*animus*) or moving away from outgroup members (which can also reflect guilt and anxiety; Ashmore, 1970).

Nonetheless, although automatic biases have been linked to negative judgments of Blacks (e.g. Jackson, 1997; Lambert, Payne, Ramsey, & Shaffer, 2005; Rudman & Lee, 2002), and female job applicants (Rudman & Glick, 2001), unambiguously harmful behaviors are seldom investigated, whether explicit or implicit bias is assessed. The present research sought to begin to fill this gap in the implicit literature. To do so, we focused on the IAT because it has borne the brunt of researchers' criticisms. In Study 1, we assessed participants' reports of their harmful actions toward Blacks in the past. Behaviors consisted of both active harm (e.g. verbal insults and physical violence) and passive harm (e.g. avoidance or exclusion). In Study 2, we measured people's willingness to cut the

budget for student organizations designed to support Jews, Asians, and Blacks. If attitude and stereotype IAT scores predict these behaviors, results would lend confidence to the IAT's ability to tap implicit prejudice, as opposed to 'mere associations' (Arkes & Tetlock, 2004, p. 268).

Study 1

Our primary aim was to examine the relationship between IAT-assessed biases and non-Black participants' self-reported harmful actions toward Blacks. To encourage candid responding, we asked participants whether they had been the victim of each behavior prior to their report of being a perpetrator. A secondary aim was to compare the attitude and stereotype IATs as predictors of harmful discrimination. To do so, we used evaluative stereotypes because they are conceptually akin to prejudice (see also Rudman et al., 2001; Wittenbrink, Judd, & Park, 1997). Because the stereotype IAT consists of both evaluative and cognitive associations, they might capture the implicit prejudice construct more completely (Breckler, 1984). Finally, we included self-report measures as a means of testing whether implicit associations can predict discrimination above and beyond explicit attitudes.

Method

Participants Sixty-four volunteers (21 male, 43 female) participated to partially fulfill an introductory psychology course requirement (M age = 20). Of these, 52 were White (81%), 6 were Asian American (9%), and 6 were Latino (9%). Data from six participants showing high error rates (> 25%) on the computer tasks were eliminated, as were data from 13 African American participants.

IAT measures The attitude IAT and stereotype IAT each used 7 White male names (e.g. John, Andrew, Peter, Brad) and 7 Black male names (e.g. Lamar, Malik, Rashan, Leroy) as the target concepts. Table 1 shows the remaining stimuli. The attitude IAT used the pleasant and unpleasant words shown in columns 1–2. The stereotype IAT used negative attributes

Table 1. Stimuli for implicit prejudice and stereotype measures (Study 1)

Attitude IAT		Stereotype IAT	
Pleasant words	Unpleasant words	Negative traits	Positive traits
sunshine	filth	lazy	ambitious
smile	death	shiftless	industrious
angel	devil	unemployed	successful
luck	slime	hostile	calm
rainbow	cancer	dangerous	trustworthy
paradise	hell	threaten	ethical
fortune	poison	violent	lawful

Note: IAT stimuli were adopted from past research (Rudman et al., 2001).

associated with Blacks (e.g. lazy, hostile) and positive attributes associated with Whites (e.g. ambitious, calm) shown in columns 3–4.

The IATs were adopted from and administered exactly as in past research (Rudman et al., 2001). The order in which participants performed the critical blocks was counterbalanced across subjects, as was the order in which participants performed the IATs (these procedural variables did not influence results). The IAT effect was computed so that high scores reflect greater tendency to associate Blacks with negative versus positive attributes, compared to when these associations were reversed. Scoring for the IAT followed recent recommendations (Greenwald, Nosek, & Banaji, 2003). Specifically, we used the D statistic because it has been shown to be less influenced by procedural variables (e.g. counterbalancing).

Explicit measures

Explicit attitudes Participants completed a feeling thermometer and the Modern Racism Scale (MRS; McConahay, 1986). The feeling thermometer asked participants to indicate, separately for African American men and White American men, the extent to which they felt positively toward each group (0 = *extremely cold*, or *unfavorable*; 99 = *extremely warm*, or *favorable*). The difference between these measures was computed such that high scores represented more positive attitudes toward White than Black men (to mirror the IAT). The MRS consists of seven

items (e.g. ‘Blacks are getting too demanding in their push for equal rights’), scored on a scale ranging from 1 (*strongly agree*) to 5 (*strongly disagree*). MRS scores were averaged ($\alpha = .82$) so that high scores reflect more anti-Black attitudes.

Discriminatory behaviors To assess harmful behaviors, we used a slightly shorter version of a past measure (Conrada et al., 2001). Participants were asked to report how often, over the course of their lifetime, they had engaged in specific actions on a scale ranging from 1 (*never*) to 7 (*very often*). The *verbal* index averaged two items pertaining to making ethnically offensive comments and jokes, either in the presence of targets or behind their backs ($r(62) = .56, p < .001$) ($M = 3.65, SD = 1.30$). The *defensive* index averaged two items pertaining to avoiding or excluding others from social gatherings and organizations because of their ethnicity ($r(62) = .69, p < .001$) ($M = 3.72, SD = 1.35$). The *offensive* index averaged three items pertaining to nonverbal hostility (e.g. giving ‘the finger’), and physically hurting targets or their property (or threatening to do so) because of their ethnicity, ($\alpha = .89; M = 2.65, SD = .80$). For each item, participants first indicated the extent to which they had been the *target* of ethnic discrimination (e.g. ‘How often have you been the target of offensive comments because of your ethnicity?’). The purpose of these items was to encourage reporting discrimination toward others (i.e. to justify participants’ own behavior). As expected, these items covaried with reports of verbal, defensive, and offensive behaviors (all r s $> .31, p$ s $< .05$).

Because they served no other purpose, they are not further discussed.¹

Procedure Volunteers participated individually in separate cubicles. Measures were administered using a computer program that presented items randomly, within each measure. To increase participants’ confidence in their anonymity, they generated their own identification number. They first completed the attitude and stereotype IATs, in counterbalanced order. They then completed the explicit attitude and behavior measures, in counterbalanced order. Consistent with past research, the order of administering the IAT and self-reports did not affect scores on either implicit or explicit measures (e.g. Greenwald et al., 2003).

Results and discussion

Preliminary analyses To assess internal consistency, we correlated practice trials with critical trials within each IAT. The coefficient for the attitude IAT was reliable ($r(62) = .69, p < .001$), as it was for the stereotype IAT ($r(62) = .71, p < .001$). Table 2 (fourth row) displays mean latencies for the implicit attitude and stereotyping measures.² On average, participants favored Whites over Blacks on both IATs, resulting in large effect sizes for both measures (attitude $d = .75$, stereotype $d = .76$).³ By contrast, negligible prejudice was reported on the MRS and the thermometer index (i.e. on average, Whites were not evaluated more favorably than were Blacks; $d = .09$).

Table 2. Summary statistics for implicit and explicit measures (Study 1)

Measure	Attitude IAT	Stereotype IAT	MRS	Thermometer index
Stereotype IAT	.33**			
MRS	.37**	.20		
Thermometer index	.25*	.16	.32*	
Mean	.43	.28	1.86	2.20
SD	.56	.37	.64	24.02

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

Note: IAT results are displayed using the *D* statistic (Greenwald et al., 2003). High scores on the IATs and the thermometer index reflect more positive evaluation of Whites compared with Blacks.

As seen in Table 2, the attitude and stereotype IATs covaried (see also Rudman et al., 2001). The two explicit measures were also related. The attitude IAT was reliably related to the MRS and to the thermometer. The stereotype IAT was positively but unreliably related to the explicit measures.

Predicting discrimination Study 1's main objective was to test whether the IAT predicts unambiguously harmful behaviors. As shown in Table 3, the attitude IAT covaried with verbal discrimination (e.g. ethnic slurs and jokes), whereas the stereotype IAT was related to each behavioral index (verbal, defensive, and offensive). Finally, the MRS reliably covaried with all three behavioral indexes, whereas the thermometer index reliably covaried with defensive and offensive, but not verbal, discrimination.

The behavioral measures were robustly related ($\alpha = .80$), and were therefore combined. A hierarchical regression analysis was then conducted to examine whether the IAT predicts unique variance in discrimination, after accounting for explicit measures. Table 4 shows the results. The discrimination index was reliably predicted by the stereotype (but not the attitude) IAT, even after accounting for the MRS and the thermometer index, which also contributed unique variance.

In sum, Study 1's focal results were the linkages shown between implicit associations and participants' history of anti-Black discrimination. The behaviors predicted by the stereotype IAT ranged from active harm (e.g. verbal slurs and personal and property violations) to more passive harm (e.g. exclusion and avoidance), whereas the attitude IAT predicted offensive comments and jokes (using bivariate analyses). Although the attitude IAT did not contribute unique variance to the discrimination index after accounting for explicit measures, this may be due to the stronger correlations between it and the direct measures, compared with the stereotype IAT. Alternatively, the stereotype IAT, because it combines beliefs with evaluation, may be a superior measure of implicit bias. For exploratory purposes, we conducted a hierarchical regression on the behavioral index in which

Table 3. Relations among intergroup orientations and discrimination (Study 1)

Behavior measure	Attitude IAT	Stereotype IAT	MRS	Thermometer
Verbal	.41**	.34**	.37**	.17
Defensive	.15	.30*	.51*	.44**
Offensive	.12	.25*	.53**	.31*

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

Note: Correlations were computed using the *D* statistic (Greenwald et al., 2003). Correlations using log transformed latencies were similar.

Table 4. Predicting harmful discrimination from explicit and implicit measures (Study 1)

Hierarchical regression model	Step	β	<i>t</i>	R^2	<i>F</i> Δ
MRS	1	.49	4.51**	.35	16.71**
Thermometer index	1	.22	2.00*		
MRS	2	.42	3.86**		
Thermometer index	2	.31	2.85**		
Attitude IAT	2	.07	.65		
Stereotype IAT	2	.33	2.98**	.44	4.53*

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

Notes: Standardized regression coefficients are shown. IAT effects used in these analyses were based on the *D* statistic (Greenwald et al., 2003).

the attitude IAT was entered first, followed by the stereotype IAT. In Step 1, the attitude IAT was a reliable predictor ($\beta = .27$, $p < .05$), but it was reduced to nonsignificance in Step 2 ($\beta = .17$, $p = .16$), suggesting that the stereotype IAT is a more effective predictor of harmful actions ($\beta = .28$, $p < .05$).

Study 2

In Study 2, we extended our analysis to include economic discrimination against Jews, Asians, and Blacks. Data concerning each group were collected over three phases, during a time period of approximately three months. Each investigation examined predictive utility for IAT-assessed evaluative stereotypes (e.g. negative attributes associated with Jews and positive

attributes associated with Christians) vis-a-vis budget cuts for minority student organizations. As in Study 1, the MRS and feeling thermometers were included for comparison purposes. The Asian-White and the Black-White phases also included an attitude IAT.

Because some investigators have argued that the IAT is a measure of environmental associations (Karpinski & Hilton, 2001; Olson & Fazio, 2004), we also included a direct measure of cultural favoritism. If IAT scores are related to perceptions that cultural stereotypes are more positive for majority than minority groups, support for this reasoning will be shown. However, Nosek and Hansen (in press), using over 50 attitude objects (including groups based on ethnicity, religion, sexual orientation, and gender) and thousands of Web site respondents, consistently found a negligible link between cultural favoritism and the IAT. Therefore, we expected to find a similar pattern.

Method

Participants All participants volunteered in exchange for partial fulfillment of their Introductory Psychology course research requirement. Only data from group members represented in the IAT were used in the analyses. In the Jewish-Christian phase, there were 89 volunteers (64 Christians, 25 Jews). Of these, 37 were men and 52 were women. In the Asian-White phase, there were 89 volunteers (59 Whites, 30 Asians). Of these, 38 were men and 51 were women. In the Black-White phase, there were 126 volunteers (89 Whites, 37 Blacks). Of these, 34 were men and 92 were women.

IAT measures

Jewish-Christian phase The stereotype IAT used six negative Jewish attributes (e.g. cheap, controlling, dominating) and six positive Christian attributes (e.g. generous, charitable, friendly). Following past research (Rudman, Greenwald, Mellott, & Schwartz, 1999), target concepts consisted of six Jewish surnames (e.g. Shapiro, Cohen, Katz) and six Christian surnames (e.g. Miller, Taylor, Johnson). The attitude IAT was not administered in this phase.

Asian-White phase Each IAT used six Asian surnames (e.g. Chang, Kwan, Yamashita) and six White names (e.g. Miller, Taylor, Johnson) as the target concepts. The attitude IAT was otherwise identical to Experiment 1's. The stereotype IAT used six negative Asian attributes (e.g. reserved, stiff, inhibited) and six positive White attributes (e.g. warm, friendly, outgoing).

Black-White phase The attitude and stereotype IATs were adopted from Experiment 1.

Explicit measures

Explicit attitudes The thermometer index was identical to Experiment 1's, with participants indicating their feelings toward the appropriate groups in each phase. A difference score was computed such that high scores indicated more positive evaluation of Christians compared with Jews, Whites compared with Asians, or Whites compared with Blacks. When necessary, the MRS was modified by replacing Blacks with either Jews or Asians as the target group (all α s > .84). The measure was scored such that high scores reflected more symbolic prejudice.

Cultural knowledge Participants were asked to rate how positive the cultural stereotypes of each group represented in the IATs were on scales ranging from 1 (*not at all*) to 10 (*extremely*). A difference score was formed so that high scores reflected judging stereotypes about majority groups as more positive than stereotypes about majority groups for each phase (e.g. Christians higher than Jews in the Jewish-Christian phase).

Economic discrimination Following past research (Haddock, Zanna, & Esses, 1993; Zanna, 2004), participants completed a budget measure that was presented as a survey conducted on behalf of the Psychology Department (i.e. separate from the main study), and was prefaced by the following statement:

We have been asked to administer this short survey as part of all of our research protocols this year, as a means of gathering student opinion. The student government has been forced to cut funding to student organizations by 20%. We ask that you help

out by recommending which organizations listed below should have their funds decreased. Current funding for each organization is listed in column 1. Place your recommended funding in column 2. Keep in mind that your suggestions should result in an approximately 20% decrease in funding. Please place this survey in the box when you are through. The results of this survey will be presented to the student government.

Eight student organizations were then listed, including the focal groups (Chabad Jewish Student Organization, Japanese Cultural Association, and Blacks United to Save Themselves) and five fillers (e.g. the PIRG organization, the drama club, and the marching band). Current funding for the focal group in each phase was listed as US\$11, 500. The difference between this and participants' recommended funding for that group was computed so that high scores indicated greater budget cuts (i.e. economic discrimination).

Procedure Upon entering the lab, participants were escorted to a separate room and asked to complete the budget recommendation measure before participating in the 'main study'. Participants placed their completed survey (subtly coded with their identification number) in a box marked 'Psychology Department Survey' to enhance the cover story. Participants were then led to a private cubicle where they performed the attitude and stereotype IATs in counterbalanced order (except in the Jewish-Christian phase, when only stereotypes were assessed), as well as the explicit measures (in the order described above). The implicit and explicit measures were administered in counterbalanced order. The IATs were administered exactly as in Experiment 1 (e.g. with task order counterbalanced). The effects of these procedural variables were non-significant in each phase. Upon completion of the measures, participants returned to the main room for a process debriefing. No subject expressed suspicion that the budget measure was part of the protocol.

Results and discussion

Preliminary analyses Internal consistency analyses of the IATs revealed reliable coefficients

for the Jewish-Christian phase ($r(87) = .61$, $p < .01$); the Asian-White phase ($r(87) = .59$ and $.60$ for the attitude and stereotype IATs, respectively, $ps < .01$); and the Black-White phase ($r(124) = .57$ and $.63$) for the attitude and stereotype IATs, respectively, $ps < .01$. Table 5 shows descriptive statistics for each phase of Study 2's data collection, as a function of group membership (Christians compared with Jews, Whites compared with Asians, and Whites compared with Blacks).

As seen in Table 5, the IATs showed the expected pattern of known groups validity, accompanied by reasonably large group difference effect sizes (all $ds > .73$), as did the explicit attitude measures (all $ds > .74$). As in past research, Jews and Asians showed reliable ingroup bias on both sets of measures (Rudman et al., 2002). Blacks demonstrated the typical pattern of showing weak implicit, but robust explicit, ingroup bias (e.g. Nosek et al., 2002). Not surprisingly, majority groups (Christians and Whites) showed greater economic discrimination, compared with minority groups (Jews, Asians, and Blacks; all $ds > .62$). Finally, the results of the cultural stereotype index (the perceived tendency for society to view majority group members more positively than minority group members) revealed that minority group members tended to report a greater discrepancy than did majority group members. However, this difference was reliable only in the Asian-White phase ($t(87) = 2.86$, $p < .01$). In sum, there was general agreement among majority and minority group members concerning how the culture viewed their groups, but little agreement concerning how the groups should be evaluated, either implicitly or explicitly.

Table 6 shows the relationships among Study 2's variables, for each phase. As expected, the IAT was unrelated to cultural stereotype index in each phase (Nosek & Hansen, in press). By contrast, the IATs and thermometer indexes were positively correlated (r s ranged from $.27$ to $.53$). The MRS covaried with the Jewish-Christian stereotype IAT and with attitude IAT scores in the Black-White phase (echoing Study 1's results). Finally, the MRS tended to be negatively linked to cultural stereotypes. That is, people who thought the

Table 5. Summary statistics for implicit and explicit measures (Study 2)

IAT measure	Attitude IAT	Stereotype IAT	MRS	Therm index	Budget index	Cultural stereotypes
<i>Jewish-Christian</i>						
Christians (<i>n</i> = 64)	–	.37	2.45	11.25	\$2480	2.35
Jews (<i>n</i> = 25)	–	–.35	1.62	–12.91	\$1200	3.04
Pooled <i>SD</i>	–	.44	.60	16.42	\$1890	2.59
Group difference <i>d</i>	–	1.64	1.38	1.47	.68	–.27
<i>Asian-White</i>						
Whites (<i>n</i> = 59)	.42	.32	2.31	7.97	\$1895	2.24
Asians (<i>n</i> = 30)	–.47	–.28	1.20	–14.17	\$632	4.07
Pooled <i>SD</i>	.63	.47	.64	17.58	\$1981	2.94
Group difference <i>d</i>	1.41	1.28	1.74	1.25	.63	–.62
<i>Black-White</i>						
Whites (<i>n</i> = 89)	.39	.28	1.70	3.18	\$1669	5.44
Blacks (<i>n</i> = 37)	.04	–.17	1.37	–16.35	\$517	6.03
Pooled <i>SD</i>	.47	.41	.44	19.61	\$1711	2.38
Group difference <i>d</i>	.74	1.09	.75	1.02	.67	–.25

Notes: IAT results are displayed using the *D* statistic (Greenwald et al., 2003). For each measure, high scores reflect greater bias against minority groups (Jews, Asians, or Blacks) or greater perceived bias in society (cultural stereotypes). The effect size (Cohen’s *d*) represents group differences. Conventional small, medium, and large effect sizes are .20, .50, and .80, respectively (Cohen, 1988).

Table 6. Correlations among implicit and explicit measures (Study 2)

Measure	IAT measures		Explicit measures		
	Attitude	Stereotype	Thermometer	MRS	Cultural stereotypes
<i>Jewish-Christian</i>					
Thermometer index	–	.53**			
Modified MRS	–	.39**	.53**		
Cultural stereotypes	–	.11	–.05	–.31**	
Budget index	–	.38*	.47**	.11	.11
<i>Asian-White</i>					
Stereotype IAT	.28**				
Thermometer index	.43**	.28**			
Modified MRS	.15	.05	.33**		
Cultural stereotypes	.02	–.07	–.08	–.12	
Budget index	.25*	.30**	.28**	.16	.10
<i>Black-White</i>					
Stereotype IAT	.47**				
Thermometer index	.42**	.27**			
MRS	.42**	.14	.30**		
Cultural stereotypes	.02	.01	–.13	–.19*	
Budget index	.23*	.18*	.08	.05	–.03

p* < .05; *p* < .01.

Note: IAT correlations were computed using the *D* statistic (Greenwald et al., 2003). Correlations using log transformed latencies were similar.

culture favored majority groups with more positive stereotypes than minority groups tended to show less symbolic prejudice (especially toward Jews and Blacks). This suggests that minority groups' justice-seeking behaviors are supported when society is perceived as biased.

Predicting discrimination Our primary aim was to examine the relationship between IAT measures and economic discrimination. Table 6 shows that the stereotype IAT was reliably linked to the budget index in each phase. That is, people who associated minority group members with negative attributes and majority group members with positive attributes were also likely to recommend budget cuts for the target minority group's student organization. The attitude IAT performed similarly in the Asian-White and Black-White phases (it was not administered in the Jewish-Christian phase). Thus, both the stereotype and attitude IAT predicted economic discrimination. The thermometer index was

reliably related to the budget measure in the Jewish-Christian and Asian-White phases, but not in the Black-White phase. By contrast, the MRS did not covary with the budget index, irrespective of the attitude object (all r s < .17, *ns*). Finally, the cultural stereotype index was not a predictor of economic discrimination (all r s < .12, *ns*).

A hierarchical regression analysis was conducted for each phase to examine whether the IAT predicts recommended budget cuts after accounting for the thermometer index (the MRS and cultural stereotypes were not included because they showed weak predictive utility). Moreover, we controlled for group identity (coded as 0 = majority, 1 = minority) to provide a more conservative test (cf. Karpinski et al., 2005). Table 7 shows the results. For the Jewish-Christian phase, discrimination was reliably predicted by group identity, the thermometer index, and the stereotype IAT. For the Asian-White phase, discrimination was marginally

Table 7. Predicting economic discrimination from explicit and implicit measures (Study 2)

Hierarchical regression model	Step	β	t	R^2	$F\Delta$
<i>Jewish-Christian</i>					
Group identity	1	-.23	2.40*	.34	11.53**
Thermometer index	1	.48	4.44**		
Group identity	2	-.27	2.95**		
Thermometer index	2	.33	2.58*		
Stereotype IAT	2	.39	3.03**	.41	9.21**
<i>Asian-White</i>					
Group identity	1	-.14	1.17	.10	4.32*
Thermometer index	1	.21	1.73		
Group identity	2	-.01	.10		
Thermometer index	2	.22	1.83		
Attitude IAT	2	.15	1.38		
Stereotype IAT	2	.27	2.41*	.17	3.61*
<i>Black-White</i>					
Group identity	1	-.24	2.38*	.09	5.93**
Thermometer index	1	.10	1.00		
Group identity	2	-.17	1.55		
Thermometer index	2	.08	.82		
Attitude IAT	2	.01	.10		
Stereotype IAT	2	.24	2.51*	.14	3.33*

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

Notes: Standardized regression coefficients are shown. IAT effects used in these analyses were based on the D statistic (Greenwald et al., 2003). Group identity was dummy coded (0 = majority, 1 = minority).

predicted by the thermometer index ($p = .07$), and reliably predicted by the stereotype IAT. For the Black-White phase, only the stereotype IAT contributed uniquely to discrimination. Thus, the stereotype IAT remained a predictor even after controlling for explicit attitudes and group identity in each phase. By contrast, the attitude IAT did not account for unique variance, and the thermometer index was only significant in the Jewish-Christian phase.

As in Study 1, we hierarchically regressed the behavioral index on the attitude IAT, followed by the stereotype IAT. For the Asian-White phase, the attitude IAT was a reliable predictor in Step 1 ($\beta = .25, p < .05$), but it was reduced to nonsignificance in Step 2 ($\beta = .16, p = .11$) after the stereotype IAT was accounted for ($\beta = .30, p < .01$). For the Black-White phase, the attitude IAT was a reliable predictor in Step 1 ($\beta = .23, p < .05$), but in Step 2 it was dramatically reduced ($\beta = .05, ns$); by contrast, the stereotype IAT was significant ($\beta = .24, p < .05$).

In sum, Study 2 showed that implicit biases predicted economic discrimination toward Jews, Asians, and Blacks, and that the stereotype IAT was either an equal or superior predictor, compared with explicit attitudes. When conditions afforded a comparison of the attitude and stereotype IATs, the latter was more effective vis-a-vis contributing unique variance. In concert with Study 1, the pattern suggests that evaluative stereotypes reflect implicit biases better than evaluative associations alone. Finally, in each phase, the cultural stereotype index was unrelated to either attitude or stereotype IATs, as in past research (Nosek & Hansen, in press). It was also unrelated to economic discrimination.

General discussion

Across two studies, the stereotype IAT predicted harmful actions toward outgroup members, even after accounting for explicit prejudice measures. In Study 1, behaviors included non-Black participants' reported history of verbal, defensive, and offensive racial discrimination. In Study 2, stereotype IATs predicted budget reductions for Jewish, Asian, and Black student organizations. The behaviors assessed ranged

in severity from racial jokes to budget cuts to blatant aggression, but each is overtly discriminatory and therefore of consequence. Moreover, we included behaviors that can be characterized as unambiguously hostile (e.g. giving the finger and physically harming an individual or their property). Although we cannot rule out the possibility that these behaviors reflected other emotions (such as fear or guilt), they undoubtedly also stem from antipathy.

Taken together, the results support conceptualizing the IAT as a measure of individual differences in automatic biases. Study 2 directly tested whether cultural favoritism influences the IAT, but found no evidence to support the hypothesis (see also Nosek & Hansen, in press). Majority and minority group members recognized the latter's lower status, but this did not influence their implicit biases. Thus, our results are not in line with the strong form of the environmental associations hypothesis, in which it is argued that IAT scores are attributable primarily to cultural, rather than personal, attitudes (Arkes & Tetlock, 2004; Olson & Fazio, 2004; Karpinski & Hilton, 2001; cf. Karpinski et al., 2005 for a more moderate version).

This is not to imply that cultural milieu has no influence on implicit biases, but rather to stress that there is no clear boundary between self and society—and this may be particularly true at the automatic level (Banaji, 2001; Devine, 1989). Indeed, there are theoretical reasons to suspect that culture can condition people's attitudes, with or without their consent (e.g. Banaji, 2001; Devine, 1989; Gaertner & Dovidio, 1986; Greenwald & Banaji, 1995). Moreover, cultural biases may be internalized for many reasons, including self-esteem, system justification, and social adjustment functions. Thus, the relationship between self and society is likely to be interdependent, even for individuals who resist being prejudiced—a fact that leads to the necessity of becoming aware of automatic biases in order to combat them.

The normativeness of implicit bias

The observation that IAT scores predict a range of discriminatory actions suggests that they are person-centered and somewhat reflective

of antipathy (see also Poehlman et al., 2004). In the same breath, we do not believe that IAT scores indicate explicit bigotry. Implicit biases are simply automatic (over-learned) evaluations; while they may reflect hostility, they also stem from many other influences (Rudman, 2004), including a natural proclivity for partisanship (Greenwald et al., 2002). In this respect, we are reminded of Allport's (1954) defense of the normality of social categorization and pre-judgment, which softened the moral sting of prejudice without removing responsibility for it (Fiske, 2004).

Indeed, a recurrent insight from response latency measures is that people tend to automatically react with preference for similar others (as they do for themselves; Greenwald & Farnham, 2000). Although this bias is condoned for many preferences (e.g. for our own children), it raises the specter of bigotry when applied to groups who do not share our genetic makeup or cultural background. Perhaps this is why some authors have argued (prematurely, in our view) that cultural biases are primarily responsible for IAT scores. But if we can view automatic biases as reflective of the normal human condition, we will be less likely to stiff-arm the messenger and hopefully, more open to becoming aware of them in order to better combat their consequences.

Limitations and future directions

In Study 1, asking participants to report past hostile behaviors was likely to evoke social desirability concerns. To counter this, we allowed participants to first report the extent to which they had suffered each behavior (Contrada et al., 2001). The goal was to afford justification for respondents' own actions and encourage honesty. Future research should compare results with and without justification items to examine their effect. If they are advantageous, the procedure may provide a template for assessing overt hostility—an important research agenda (Mackie & Smith, 1998).

In two studies, the stereotype IAT predicted a range of discriminatory behaviors, in support of its construct validity. This was true when we controlled for explicit attitudes in both studies, and for group identity in Study 2. The relatively

weak performance of the attitude IAT (after accounting for the stereotype IAT, as well as explicit measures) suggests that evaluative beliefs may capture the implicit prejudice construct better than simple good-bad associations. This advantage may stem from the fact that negative outgroup stereotypes (and positive ingroup stereotypes) afford more justification for discrimination than associating ingroup and outgroup members with pleasant and unpleasant words. However, considerably more research that affords a comparison between implicit attitudes and stereotypes is needed before we can have confidence in the stereotype IAT's superior predictive utility.

In addition, it would be interesting to uncover moderators of reactions to IAT scores. To date, there are indications that people who are motivated not to appear racist, or who are anxious about their scores, tend to react defensively (Frantz, Cuddy, Burnett, Ray, & Hart, 2004; Monteith, Voils, & Ashburn-Nardo, 2001). But some people respond to their IAT scores with greater equanimity. For example, the first author demonstrates her automatic biases in the classroom, to create an atmosphere of trust. Placed in this context, implicit orientations become more normative and less threatening—not because society is to be blamed for them, but because growing up in a culture where some people are valued more than others is likely to permeate our private orientations, no matter how discomfiting the fact (Banaji, 2001; Devine, 1989). For this reason, the IAT is a powerful educational tool, as it opens people up to discussions about social justice that might otherwise be dismissed as antiquated (Bombardieri, 2005). Even people who argue that biases can be rational (e.g. Arkes & Tetlock, 2004) can appreciate the disconcerting fact that Blacks as a group are automatically associated with negative attributes for many Whites.

Future research should also continue to compare the predictive utility of implicit and explicit attitudes. Although it has been argued that implicit biases best predict spontaneous behaviors (Fazio & Olson, 2003), self-reports (a controlled behavior) often covary with them. In fact, validity for response latency measures has often relied

on controlled judgments (e.g. Livingston, 2002; Rudman & Lee, 2002). Although some implicit measures appear to be better at predicting automatic versus controlled behavior (Dovidio et al., 2002), in other cases the behaviors have been a mix of automatic and controlled actions (Fazio et al., 1995; McConnell & Leibold, 2001). The present research used behavioral measures that were likely more controlled than automatic (reporting past discrimination in Study 1 and recommending budget cuts in Study 2). Thus, there is no clean, process-driven divide by which to define the predictive utility of implicit and explicit responses.

Finally, the link between explicit attitudes and behavior is moderated by numerous factors; it cannot be otherwise for implicit attitudes (Fazio & Olson, 2003; Karpinski et al., 2005; Nosek, 2005). Moreover, behavior directed toward individuals may be dramatically different from group-related actions; our brains react differently when we individuate (Wheeler & Fiske, 2005), and so do we (Brewer, 1988; Fiske & Neuberg, 1990). Thus, we caution against assuming that IAT scores reflect how people will generally behave toward outgroup members.

Notes

1. Readers might be interested in the relationships among being victimized and the implicit and explicit measures. Participants who reported being excluded by outgroup members tended to show high scores on the stereotype IAT ($r(62) = .32, p < .01$), and marginally, on the MRS and the thermometer (both $r(62) = .22, ps < .08$). Targets of physical (or threatened) aggression, and those who received the finger, also scored higher on the stereotype IAT (both $r(62) = .24, ps = .05$). No other relationships approached significance (all $r(62) < .15, ns$).
2. The data for each trial block included mean response latency (in ms) and error rates. Response latencies greater than 3000 ms and less than 300 ms were recoded as 3000 and 300 ms, respectively; the first two trials of each block were dropped because of their typically lengthened latencies. Analyses of participants' accuracy revealed low error rates on critical trials (an average of 5%); error trials were included in IAT effect scores. Following Greenwald et al. (2003), we computed the D statistic for use in correlational analyses. Results using log transformed IAT scores were similar.
3. By convention, small, medium, and large effect sizes correspond to .20, .50, and .80, respectively (Cohen, 1988).

Acknowledgments

This research was partially supported by Grants BCS-0109997 and BCS-0417335 from the National Science Foundation. Preparation of this article was partially supported by Grant BCS-0417335 from the National Science Foundation to the first author.

References

- Allport, G. W. (1954). *The nature of prejudice*. New York: Addison-Wesley.
- Arkes, H. R., & Tetlock, P. E. (2004). Attributions of implicit prejudice, or 'would Jesse Jackson "fail" the Implicit Association Test?' *Psychological Inquiry, 15*, 257–278.
- 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.
- Ashmore, R. D. (1970). The problem of intergroup prejudice. In B. E. Collins (Ed.), *Social psychology: Social influence, attitude change, group processes, and prejudice* (pp. 245–339). Reading, MA: Addison-Wesley.
- Banaji, M. R. (2001). Implicit attitudes can be measured. In H. L. Roediger, J. S. Nairne, I. Neath, & A. Surprenant (Eds.), *The nature of remembering: Essays in honor of Robert G. Crowder* (pp. 117–150). Washington D.C.: American Psychological Association.
- Blair, I. V. (2001). Implicit stereotypes and prejudice. In G. B. Moskowitz (Ed.), *Cognitive social psychology: The Princeton symposium on the legacy and future of social cognition* (pp. 359–374). Mahwah, NJ: Erlbaum.
- Blair, I. V. (2002). The malleability of automatic stereotypes and prejudice. *Personality and Social Psychology Review, 6*, 242–261.
- Bombardieri, M. (2005, April 8). Summers displays new understanding of women's careers. *The Boston Globe*, p. B3.
- Breckler, S. J. (1984). Empirical validation of affect, behavior, and cognition as distinct components of attitude. *Journal of Personality and Social Psychology, 6*, 1191–1205.

- Brewer, M. B. (1988). A dual-process model of impression formation. In T. K. Srull & R. W. Wyer, Jr. (Eds.), *Advances in social cognition* (Vol. 1, pp. 1–36). Hillsdale, NJ: Erlbaum.
- Cohen, J. (1988). *Statistical power for the behavioral sciences*. Hillsdale, NJ: Erlbaum.
- Contrada, R. J., Ashmore, R. D., Gary, M. L., Coups, E., Egeth, J. D., Sewell, A. et al. (2001). Measures of ethnicity-related stress: Psychometric properties, ethnic group differences, and associations with well-being. *Journal of Applied Social Psychology, 31*, 1775–1820.
- Cunningham, W. A., Preacher, K. J., & Banaji, M. R. (2001). Implicit attitude measures: Consistency, stability, and convergent validity. *Psychological Science, 12*, 163–170.
- 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., & Gaertner, S. L. (2002). Implicit and explicit prejudice and interracial interactions. *Journal of Personality and Social Psychology, 82*, 62–68.
- 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.
- Fiske, S. T. (2004). Intent and ordinary bias: Unintended thought and social motivation create casual prejudice. *Social Justice Research, 17*, 117–127.
- Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 23, pp. 1–74). New York: Academic Press.
- Frantz, C. M., Cuddy, A. J. C., Burnett, M., Ray, H., & Hart, A. (2004). A threat in the computer: The race Implicit Association Test as a stereotype threat experience. *Personality and Social Psychology Bulletin, 30*, 1611–1624.
- Gaertner, S. L., & Dovidio, J. F. (1986). The aversive form of racism. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 61–90). Orlando, FL: Academic Press.
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review, 102*, 4–27.
- Greenwald, A. G., & Farnham, S. D. (2000). Using the Implicit Association Test to measure self-esteem and self-concept. *Journal of Personality and Social Psychology, 79*, 1022–1038.
- Greenwald, A. G., McGhee, D., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology, 74*, 1464–1480.
- Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology, 85*, 197–216.
- Greenwald, A. G., Pickrell, J. E., & Farnham, S. D. (2002). Implicit partisanship: Taking sides for no reason. *Journal of Personality and Social Psychology, 83*, 367–379.
- Haddock, G., Zanna, M. P., & Esses, V. M. (1993). Assessing the structure of prejudicial attitudes: The case of attitudes toward homosexuals. *Journal of Personality and Social Psychology, 24*, 189–205.
- Islam, M. R., & Hewstone, M. (1993). Dimensions of contact as predictors of intergroup anxiety, perceived out-group variability, and out-group attitude: An integrative model. *Personality and Social Psychology Bulletin, 19*, 700–710.
- Jackson, J. R. (1997). *Automatically activated racial attitudes*. Unpublished doctoral dissertation, Indiana University.
- Jost, J. T., Pelham, B. W., & Carvallo, M. R. (2002). Non-conscious forms of system justification: Implicit and behavioral preferences for higher status groups. *Journal of Experimental and Social Psychology, 38*, 586–602.
- Karpinski, A., & Hilton, J. L. (2001). Attitudes and the Implicit Association Test. *Journal of Personality and Social Psychology, 81*, 774–788.
- Karpinski, A., Steinman, R. B., & Hilton, J. L. (2005). Attitude importance as a moderator of the relationship between implicit and explicit attitude measures. *Personality and Social Psychology Bulletin, 31*, 949–962.
- Lambert, A. J., Payne, B. K., Ramsey, S., & Shaffer, L. M. (2005). On the predictive validity of implicit attitude measures: The moderating effect of perceived group variability. *Journal of Experimental Social Psychology, 41*, 114–128.
- 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.

- Mackie, D. M., & Smith, E. R. (1998). Intergroup relations: Insights from a theoretically integrative approach. *Psychological Review*, 105, 499–529.
- 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. 91–126). Orlando, FL: Academic Press.
- McConnell, A. R., & Leibold, J. M. (2001). Relations among the Implicit Association Test, explicit attitudes, and discriminatory behavior. *Journal of Experimental Social Psychology*, 37, 435–442.
- Monteith, M. J., Voils, C. I., & Ashburn-Nardo, L. (2001). Taking a look underground: Detecting, interpreting, and reacting to implicit biases. *Social Cognition*, 19, 395–417.
- Nosek, B. A. (2005). Moderators of the relationship between implicit and explicit evaluation. *Journal of Experimental Psychology: General*, 134, 565–584.
- Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002). Harvesting implicit group attitudes and beliefs from a demonstration web site. *Group Dynamics*, 6, 101–115.
- Nosek, B. A., & Hansen, J. J. (in press). The associations in our heads belong to us: Searching for attitudes and knowledge in implicit evaluation. *Cognition and Emotion*.
- Olson, M. A., & Fazio, R. H. (2004). Reducing the influence of extra-personal associations on the Implicit Association Test: Personalizing the IAT. *Journal of Personality and Social Psychology*, 86, 653–667.
- Poehlman, T. A., Uhlmann, E., Greenwald, A. G., & Banaji, M. R. (2004). *Measuring and using the Implicit Association Test: III. Meta-analysis of predictive validity*. Manuscript submitted for publication.
- Rudman, L. A. (2004). Sources of implicit attitudes. *Current Directions in Psychological Science*, 13(2), 79–83.
- Rudman, L. A., Ashmore, R. D., & Gary, M. L. (2001). ‘Unlearning’ automatic biases: The malleability of implicit stereotypes and prejudice. *Journal of Personality and Social Psychology*, 81, 856–868.
- Rudman, L. A., Feinberg, J. M., & Fairchild, K. (2002). Minority members’ implicit attitudes: 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.
- Rudman, L. A., Greenwald, A. G., Mellott, D. S., & Schwartz, J. L. K. (1999). Measuring the automatic components of prejudice: Flexibility and generality of the Implicit Association Test. *Social Cognition*, 17(4), 1–29.
- Rudman, L. A., & Lee, M. R. (2002). Implicit and explicit consequences of exposure to violent and misogynous rap music. *Group Processes & Intergroup Relations*, 5, 133–150.
- Sears, D. (1988). Symbolic racism. In P. Katz & D. Taylor (Eds.), *Eliminating racism: Profiles in controversy* (pp. 53–84). New York: Plenum.
- Son Hing, L. S., Li, W., & Zanna, M. P. (2002). Inducing hypocrisy to reduce prejudicial responses among aversive racists. *Journal of Experimental Social Psychology*, 38, 71–78.
- Stephan, W. G., & Stephan, C. W. (1985). Intergroup anxiety. *Journal of Social Issues*, 41, 157–175.
- Uhlmann, E., & Poehlman, T. A. (2005). *Automatic associations: Personal attitudes or cultural knowledge?* Unpublished manuscript, Yale University.
- Wheeler, M. E., & Fiske, S. T. (2005). Controlling racial prejudice: Social cognitive goals affect amygdala and stereotype activation. *Psychological Science*, 16, 56–63.
- 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.
- Zanna, M. P. (2004). The naive epistemology of a working social psychologist (or the working epistemology of a naive social psychologist): The value of taking ‘temporary givens’ seriously. *Personality and Social Psychology Review*, 8, 210–218.

Paper received 17 June 2005; revised version accepted 6 June 2006.

Biographical notes

LAURIE RUDMAN is an associate professor in the Psychology Department at Rutgers University. Her primary research interests are prejudice, stereotypes, and implicit social cognition.

RICHARD ASHMORE is a professor in the Psychology Department at Rutgers University. His primary research interests are self and identity and intergroup relations.