Journal of Experimental Social Psychology **36,** 316–328 (2000)
doi:10.1006/jesp.1999.1418, available online at http://www.idealibrary.com on IDEAL®

Automatic Preference for White Americans: Eliminating the Familiarity Explanation

Nilanjana Dasgupta

New School University

Debbie E. McGhee and Anthony G. Greenwald

University of Washington

and

Mahzarin R. Banaji

Yale University

Received March 19, 1999; revised September 30, 1999; accepted November 16, 1999

Using the Implicit Association Test (IAT), recent experiments have demonstrated a strong and automatic positive evaluation of White Americans and a relatively negative evaluation of African Americans. Interpretations of this finding as revealing pro-White attitudes rest critically on tests of alternative interpretations, the most obvious one being perceivers' greater familiarity with stimuli representing White Americans. The reported experiment demonstrated that positive attributes were more strongly associated with White than Black Americans even when (a) pictures of equally unfamiliar Black and White individuals were used as stimuli and (b) differences in stimulus familiarity were statistically controlled. This experiment indicates that automatic race associations captured by the IAT are not compromised by stimulus familiarity, which in turn strengthens the conclusion that the IAT measures automatic evaluative associations. © 2000 Academic Press

National surveys indicate that racism in American society has declined steadily over the past 50 years (Schuman, Steeh, & Bobo, 1997). Despite this

This research was supported by grants from the National Science Foundation (SBR-9422242, SBR-9710172, SBR-9422241, and SBR-9709924) and from the National Institute of Mental Health (MH-41328, MH-01533, and MH-57672). The authors thank William Cunningham and Scott Ottaway for their comments on an earlier draft.

Address correspondence and reprint requests to Nilanjana Dasgupta, Department of Psychology, New School University, Graduate Faculty, 65 Fifth Avenue, New York, NY 10003, e-mail: dasguptn@newschool.edu; Debbie E. McGhee and Anthony G. Greenwald, Department of Psychology, University of Washington, P.O. Box 351525, Seattle, WA 98195-1525, e-mail: demcghee@u.washington.edu or agg@u.washington.edu; or Mahzarin R. Banaji, Department of Psychology, Yale University, P.O. Box 208205, New Haven, CT 06520-8205, e-mail: mahzarin.banaji@yale.edu.

0022-1031/00 $35.00

optimistic finding, other research using indirect measures suggests that subtle and implicit forms of prejudice and discrimination remain pervasive (Crosby, Bromley, & Saxe, 1980; Fiske, 1998). One explanation for the discrepancy between explicit and implicit attitudes is that explicit attitudes assessed by self-report measures may be more susceptible to self-presentation bias, whereas implicit attitudes captured by indirect measures may be less vulnerable to such influences. Even when self-presentation does not threaten the validity of the data, subtle forms of stereotypes and prejudice may be expressed unconsciously without perceivers' awareness or control (Banaji & Greenwald, 1995; Banaji, Hardin, & Rothman, 1993; Blair & Banaji, 1996; Fazio, Jackson, Dunton, & Williams, 1995; Greenwald & Banaji, 1995; Greenwald, McGhee, & Schwartz, 1998; Wittenbrink, Judd, & Park, 1997; also see the *Journal of Experimental Social Psychology*, **33,** special issue on unconscious prejudice and stereotyping).

Substantial progress has been made in uncovering implicit and automatic prejudice with the introduction of various indirect measures including the recently developed Implicit Association Test (IAT; Greenwald et al., 1998). The IAT uses response latency to assess the relative strength with which attitude objects are associated with particular evaluations. In the IAT, subjects classify stimuli representing racial groups (e.g., Black and White faces) and evaluative attributes (e.g., pleasant and unpleasant words) using two designated keys. Subjects typically perform this task more quickly and easily when pleasant attributes share the same response key with White than Black stimuli and unpleasant attributes share the same key with Black than White stimuli (Greenwald et al., 1998; Ottaway, Hayden, & Oakes, in press). Even subjects who are told that the IAT measures undesirable racist attitudes and who explicitly self-report egalitarian attitudes find it difficult to control their biased responses (Kim & Greenwald, 1998). In other words, IAT responses are considered automatic because they are expressed without intention or control (cf. Bargh, 1994; Greenwald & Banaji, 1995), although perceivers may become aware of the attitude object under scrutiny during the task.

Several researchers have demonstrated automatic race preference using the IAT (Greenwald et al., 1998; Ottaway et al., in press). However, the interpretation that IAT effects reveal automatic White preference rests critically on tests that rule out alternative explanations. Focusing on a potential measurement confound, one explanation argues that automatic race evaluations captured by this task may not reflect genuine White preference, but rather participants' greater familiarity with White names (e.g., Wendy, Brad) compared to Black names (e.g., Latisha, Malik) used in the task.[1] This explanation is consistent with findings showing that frequent exposure to (and hence increased familiarity with) previously neutral stimuli enhances self-reported preference for those stimuli (Zajonc, 1968). If the IAT is to be regarded as a valid measure of automatic attitudes, it must first disentangle attitudes from the effects of stimulus familiar-

[1] This argument has been repeatedly raised whenever the authors have presented data on automatic race attitudes.

ity. Such efforts are likely to have broad impact because of the growing usage of this measure as an indirect indicator of attitude.

Recently, Ottaway and colleagues examined the role of stimulus familiarity on automatic race attitudes using stimuli (Black, Hispanic, and White names) equated on self-reported familiarity and objective frequency. They found that White subjects expressed automatic preference for White over Black and Hispanic Americans even though name familiarity and frequency were controlled. The present experiment seeks to extend Ottaway et al.'s research by providing two additional and independent tests assessing the influence of stimulus familiarity on IAT responses.

In one IAT, racial groups were represented by pictures of nonfamous Black and White individuals instead of racially identifiable names previously used by Greenwald et al. (1998) and Ottaway et al. (in press). Whereas Black and White names maybe differentially familiar to subjects, pictures of nonfamous Black and White Americans ought to be equally unfamiliar to all, thereby controlling for the effect of stimulus familiarity on automatic responses.

A second test of the role of stimulus familiarity was provided by another IAT in which name stimuli were used. Subjects' familiarity with these names was assessed separately and a statistical technique used to determine the magnitude of White preference when Black and White names were equally familiar. Taken together, the two IATs provide a unique contribution by measuring and manipulating stimulus familiarity with two techniques not used in previous research and, additionally, by examining the stability of race preference effects across different types of stimuli.

We predicted that subjects would associate positive attributes more quickly and easily with White than Black exemplars regardless of their familiarity with those exemplars. If obtained, the interpretation that these associations reveal automatic race attitudes can be offered on more secure empirical grounds. The potential influence of familiarity on automatic associations is also of broader interest because it may have implications for the measurement of attitudes using the IAT in other domains where familiarity may be a confound (e.g., academic and consumer preferences and attitudes toward other groups).

EXPERIMENT

Familiarity with Black- and White-American names was assessed first followed by two IATs that measured subjects' automatic race evaluations using picture and name stimuli. Name familiarity was operationalized as the speed and accuracy with which subjects recognized and differentiated real names (both Black and White) from pseudonames that served as control stimuli. Two tasks provided independent judgments of familiarity for each group (differentiation of Black vs pseudonames and White vs pseudonames).

If subjects are familiar with real names, they should be less likely to confuse them with pseudonames, thus producing fast reaction times on the real name–pseudoname discrimination task. Because White names tend to be more familiar than Black names, differentiation of White names from pseudonames was pre-

Sequence	1	2	3	4	5
Task Description	Target concept Discrimination	Attribute Discrimination	Target+attribute Combined task	Reversed Target discrimination	Reversed target+attribute combined task
Task instructions	• BLACK WHITE •	• unpleasant pleasant •	• BLACK • unpleasant pleasant • WHITE •	BLACK • • WHITE	BLACK • • unpleasant pleasant • • WHITE
Sample	JOSH •	Gentle •	• MALIK	• JOSH	• BRANDON
stimuli	• LAMAR	• bomb	Joy •	LAMAR •	gentle •
	• MALIK	• death	JUSTIN •	MALIK •	LAMAR •
	ANDREW •	Joy •	• Death	• ANDREW	• grief
	BRANDON •	• disaster	• RASAAN	• BRANDON	pleasure •
	• JAMEL	• poison	• poison	JAMEL •	MALIK •
	• RASAAN	Pleasure •	Paradise •	RASAAN •	• tragedy
	JUSTIN •	Paradise •	ANDREW •	• JUSTIN	• JUSTIN

FIG. 1. Schematic description of the Implicit Association Test. The IAT procedure involved a series of five steps (numbered columns). The first two steps introduce target concepts (Black vs White) and attributes (pleasant vs unpleasant), each of which is assigned to a left or right response indicated by black circles. The third step combines the target and attribute classifications such that particular types of targets and attributes are mapped onto the same response key (e.g., White + pleasant in Step 3). The fifth step recombines the target + attribute classifications after reversing target responses in Step 4. Automatic preference for White Americans (i.e., the IAT effect) = (RT for Black + pleasant and White + unpleasant combinations) minus (RT for Black + unpleasant and White + pleasant combinations).

dicted to be more quick and accurate than differentiation of Black names from pseudonames. The difference in judgment speed between White name–pseudoname and Black name–pseudoname tasks provided a measure of relative familiarity with White versus Black names.

Design of the IAT

In the IAT, subjects classify target concepts (represented by Black and White exemplars) and attributes (represented by pleasant and unpleasant words) using two designated keys (see Fig. 1). Subjects typically perform this task quickly and easily when White and pleasant stimuli share the same response key or Black and unpleasant stimuli share the same key (abbreviated as White + pleasant and Black + unpleasant respectively; Step 3 of Fig. 1). In contrast, they perform the task more slowly and with greater difficulty for opposite combinations of stimuli (Black + pleasant and White + unpleasant; Step 5 of Fig. 1). The difference in response latencies for pro-White (White + pleasant and Black + unpleasant) versus pro-Black trials (Black + pleasant and White + unpleasant) provides a measure of relative automatic preference for European- compared to African-Americans.[2] The order of Steps 3 and 5 is counterbalanced across subjects.

[2] At present the IAT technique is unable to separate White + pleasant from Black + unpleasant associations, although this issue is under investigation in other research. For now, latencies for White + pleasant and Black + unpleasant associations are combined and compared with Black + pleasant and White + unpleasant associations. The data are interpreted as a measure of relative automatic preference for one group over the other.

Subjects

Seventy-five students (41 females, 34 males; 35 Caucasian, 22 Asian, 18 other ethnicities) enrolled in introductory psychology courses at the University of Washington participated in exchange for extra course credit.

Procedure

Explicit attitude measures. Five questionnaires assessing explicit race attitudes were administered first. These included: feeling thermometers, semantic differential scales, Modern Racism Scale (McConahay, Hardee, & Batts, 1981), and Diversity and Discrimination scales (Wittenbrink et al., 1997). Detailed descriptions of all explicit measures are provided by Greenwald et al. (1998).

Name recognition. Next, two name-recognition tasks were administered on desktop computers equipped with Windows 95 operating systems.[3] These tasks, involving classification of White names vs pseudonames and Black names vs pseudonames, were completed in counterbalanced order. Real names were represented by eight typically White names (e.g., Wendy, Brandon) and eight typically Black names (e.g., Latoya, Malik) with an equal number of males and females in each category. In addition, 16 pseudonames (e.g., Lesir, Lucena) were created using two criteria: (a) all were easily pronounceable, and (b) none were obviously similar to real Black and White names. Moreover, none of the pseudonames appeared in the U.S. census, lending credence to the claim that they were false names.[4]

Subjects were instructed to classify stimulus names as known versus unknown for a 16-trial practice block followed by a 32-trial data-collection block. Stimuli were presented sequentially and subjects' keypress response initiated the next trial. Incorrect classifications were followed by error feedback (the word "ERROR"). Different sets of pseudonames were used in the two tasks and counterbalanced across subjects.

IAT. Picture and name IATs were administered after explicit measures and name-recognition tasks had been completed. In the name IAT, race categories were represented by the Black and White names described above. The picture IAT used eight Black and eight White facial photographs acquired from the home page of a public university. Pictures maximized uniformity of dress (within sex) and facial expression (all targets were smiling). All pictures were 104 pixels wide, 132 pixels tall, and in 256-color grayscale format. Evaluative attributes included 16 words, 8 pleasant and 8 unpleasant (e.g., gentle, happy, disaster, and grief), selected from Bellezza, Greenwald, and Banaji (1986).

In the IAT, stimuli appeared within a centered white window against a light-blue background. Names appeared in black uppercase letters and evaluative

[3] The program was written using Inquisit created by Sean C. Draine, published by Millisecond Software.

[4] The remaining pseudonames were: Corris, Arton, Julren, Nekar, Yalton, Anadri, Birana, Lystua, Sirris, Andon, Turlen, Yaslon, Arisly, and Bralla.

attributes in black lowercase letters. Reminder labels were positioned above the stimuli on the left and right side. These reminders read "WHITE" and "BLACK" for single target-classification blocks, "pleasant" and "unpleasant" for single attribute-classification blocks. Mixed target + attribute blocks were also accompanied by appropriate labels (e.g., "Pleasant or BLACK" and "Unpleasant or WHITE"). Incorrect classifications were followed by error feedback (i.e., the word "ERROR"). Summary feedback was provided at the end of each block informing subjects about their average response latency and percentage of errors for that block.

All practice tasks in the IAT were administered in five blocks of 32 trials each. Data-collection tasks consisting of combined target + attribute classifications were administered in four blocks of 33 trials each. Within each block, stimuli were randomly selected without replacement and no more than two consecutively presented stimuli belonged to the same category. Further details about the IAT procedure are provided by Greenwald et al. (1998).

Each picture and name IAT included two within-subjects factors: (a) IAT combinations (Black + pleasant vs White + pleasant) and (b) Type of IAT (picture vs name). In addition, two between-subjects counterbalancing factors were also included: (a) Order of IAT combinations (Black + pleasant first vs White + pleasant first) and (b) Order of IATs (name IAT first vs picture IAT first).

Results

Data preparation. To correct for anticipatory responses and momentary inattention, latencies less than 300 ms and greater than 3000 ms were recoded as 300 and 3000 ms respectively for name-recognition tasks and IATs. Latencies were log transformed to normalize the distribution.

IAT. IAT effects were computed separately for name and picture tasks by subtracting the mean latency for White + pleasant and Black + unpleasant blocks (pro-White stimulus combinations) from the mean latency for Black + pleasant and White + unpleasant blocks (pro-Black stimulus combinations). Thus, positive difference scores indicate stronger pleasant associations with White compared to Black and stronger unpleasant associations with Black compared to White.

We predicted that both IATs would reveal substantial automatic preference for White Americans such that pleasant attributes would be associated more quickly and easily with White than Black stimuli and unpleasant attributes with Black than White stimuli. As predicted (see Fig. 2), both picture and name IATs revealed significantly faster response latencies when White exemplars were associated with pleasant attributes and Black with unpleasant than for opposite combinations of stimuli (average IAT effect = 111 ms; $d = .93$; $F(1, 63) = 96.45$, $p = 10^{-14}$). The magnitude of the White preference effect was significantly larger for the name IAT than the picture IAT (IAT effects = 140 vs 81

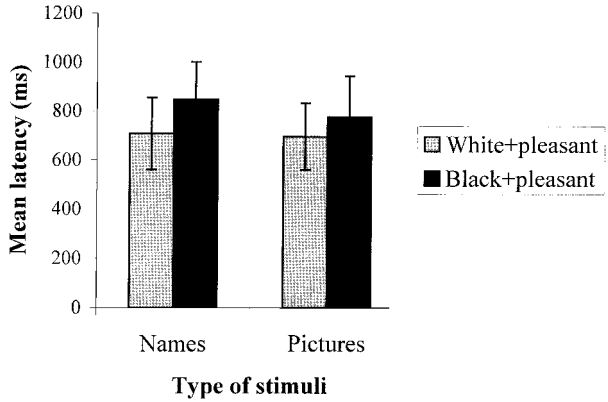

FIG. 2. Automatic associations (in milliseconds): IAT combinations by Type of stimuli. IAT effect (in milliseconds) = (Black + pleasant and White + unpleasant) minus (White + pleasant and Black + unpleasant) combinations. Name IAT effect = 140 ms (d = .93) and Picture IAT effect = 81 ms (d = .53). The analysis of variance was conducted using log-transformed response latencies. However, untransformed response latencies are presented to help interpretations.

ms; ds = .93 vs .53; $F(1, 63) = 17.91$, $p = 10^{-5}$).[5] Taken together, these findings demonstrated automatic White preference using both picture and name stimuli. Furthermore, the data revealed race bias even when pictures of unfamiliar Black and White individuals were used as stimuli, demonstrating that differential stimulus familiarity did not produce the obtained race bias effect.

Name recognition. A relative name-recognition score was computed for each subject by subtracting the mean latency (or error rate) in the White name–pseudoname condition from that in the Black name–pseudoname condition. As expected, subjects were significantly faster at differentiating White names from pseudonames than Black names from pseudonames (mean latencies = 711 vs 928 ms; difference = 217 ms; $F(1, 71) = 279.66$, $p = 10^{-26}$). They were also more accurate at distinguishing White names from pseudonames than Black names from pseudonames (mean errors = 3% vs 18%; difference = 15%; $F(1, 71) = 110.95$, $p = 10^{-16}$).

Additional analyses were conducted to buttress the argument that slower responses in the Black name–pseudoname than White name–pseudoname task were produced by subjects' unfamiliarity with real Black names rather than greater objective similarity between Black names and pseudonames than between White names and pseudonames. If stimulus *similarity* was responsible for slow latencies in the Black–pseudoname task, then those latencies ought *not* to be correlated with another task in which Black names were contrasted with clearly dissimilar stimuli (e.g., the Black–White practice task in the IAT). Following the same logic, if stimulus *dissimilarity* had produced faster latencies in the White name–pseudoname task, then those latencies *ought* to be correlated with another

[5] Also, response latencies were faster in the picture than name IAT, $F(1, 63) = 27.42$, $p = 10^{-6}$.

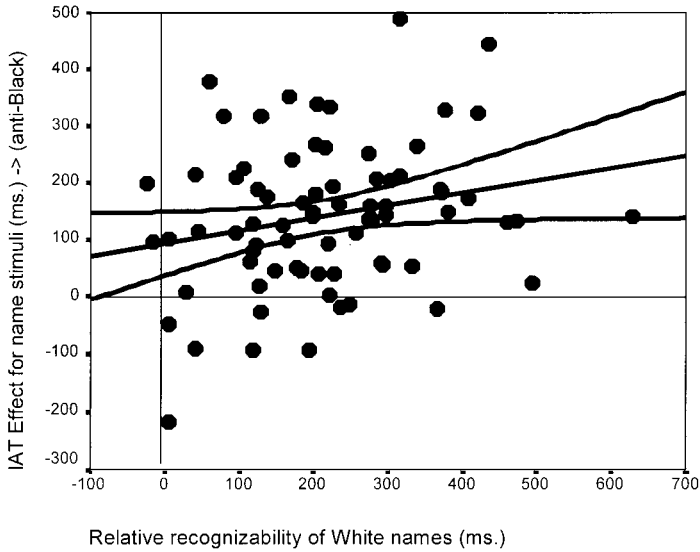

FIG. 3. The relationship between relative recognizability of racial names and automatic associations (IAT effect). The *y* intercept (i.e., IAT effect) = 92 ms; $t = 2.69$, $p < .009$; slope = .27, $t = 2.36$, $p = .02$. For name recognizability, larger numbers indicate faster recognition of White compared to Black names. This regression analysis was conducted using log-transformed response latencies. However, untransformed response latencies are presented to facilitate interpretation. Regression line is bounded by 95% confidence intervals.

task in which White names were contrasted with clearly dissimilar stimuli (e.g., Black–White practice IAT task). However, partial correlations revealed that the Black–White IAT task was significantly correlated with both Black name–pseudoname ($r = .32$, $p = .005$) and White name–pseudoname tasks ($r = .53$, $p < .0009$) even after controlling for general individual differences in responding as assessed by the pleasant–unpleasant practice task. These correlations do not support the similarity alternative.

Regression analyses. The relationship between name familiarity and automatic race preference was examined using a regression analysis technique similar to one developed by Greenwald, Klinger, and Schuh (1995). If automatic White preference was produced by greater familiarity with White names, then the magnitude of the IAT effect should be nonsignificant when White and Black names are equally familiar. However, if race preference was not completely produced by greater familiarity with White names, then the magnitude of the IAT effect should remain significant when both types of names are equally familiar.

In the regression analysis, relative recognition of White versus Black names was used as an independent variable to predict the magnitude of the IAT effect. In Fig. 3, equal familiarity with Black and White names is represented by *x* coordinate = 0 and the magnitude of the IAT effect for equally familiar race stimuli is represented by the *y* intercept of the regression line. The regression

TABLE 1
Correlations between Explicit and Implicit Attitude Measures

	Name Recognition	Name IAT	Picture IAT	Semantic Differential	Feeling Therm.	Modern Racism	Discrimination Scale
Name recognition	1.00						
IAT Effect, names	0.23*	1.00					
IAT Effect, pictures	0.18	0.39*	1.00				
Semantic Differential	−0.06	*0.24**	*0.38**	1.00			
Feeling Thermometer	−0.04	*0.23**	*0.21*	**0.58***	1.00		
Modern Racism Scale	−0.12	*−0.13*	*0.02*	**0.44***	**0.33***	1.00	
Discrimination Scale	−0.11	*0.01*	*0.08*	**0.45***	**0.36***	**0.71***	1.00
Diversity Scale	−0.12	*−0.02*	*0.14*	**0.45***	**0.39***	**0.63***	**0.71***

Note. Significant correlations are indicated with asterisks for $r = .23$, $p = .05$; $r = .295$, $p = .01$; and $r = .325$, $p = .005$. Entries in bold indicate correlations among explicit measures. Entries in italics indicate correlations between implicit and explicit measures. For name recognition, higher numbers indicate faster recognition of White compared to Black names. All other measures were coded such that higher numbers indicate greater preference for White over Black.

technique also examined whether the IAT effect changed systematically as White names became more recognizable than Black names (i.e., when $x > 0$), which is represented by the slope of the regression line.

As shown in Fig. 3, a significant IAT effect was obtained when Black and White names were equally recognizable. This finding is indicated by a significantly positive y intercept (y intercept = 92 ms, $t = 2.69$, $p = .009$). Moreover, increased recognizability of White names compared to Black names produced a modest corresponding increase in the magnitude of the IAT effect (slope = .27, $t = 2.36$, $p = .02$).[6]

Explicit measures. For semantic differential and thermometer measures, difference scores were computed by subtracting ratings of Black from White Americans, such that higher scores indicated relative preference for White over Black. Table 1 presents correlations between explicit and implicit measures. The five explicit measures formed a cluster accounting for all correlations greater than .50 (average $r = .51$). By contrast, the average correlation between explicit and implicit measures was $r = .12$. Specifically, correlations between explicit attitude and belief questionnaires (MRS, Diversity and Discrimination scales) and the IATs were nonsignificant (average $r = .02$). However, correlations among feeling thermometers, semantic differential scales, and IATs (all of which

[6] For two reasons, the Black–White practice task from the IAT is not an appropriate substitute for the real-name–pseudoname discrimination task in the regression analysis. First, the Black–White practice task cannot statistically represent a scenario in which Black and White names are equally familiar because subjects' reaction times on this task have no absolute zero point. Second, the Black–White practice task also does not provide an accurate test of how the IAT effect changes (typically captured by the slope of the regression line) when Black names are perceived to be less familiar than White names. This is because the speed with which subjects recognize Black names in that practice task is inherently confounded with the speed with which they recognize White names. By contrast, the real-name–pseudoname discrimination tasks provide a better indirect measure of name familiarity because they are not susceptible to the two problems mentioned above.

assessed differential evaluations of race) were small but significant (average r = .27). The overall correlation between implicit and explicit attitude measures (average r = .12) is consistent with previous research (e.g., Devine, 1989; Dovidio et al., 1997; Fazio et al., 1995; Greenwald et al., 1998; but see Lepore & Brown, 1997; Wittenbrink et al., 1997). For instance, Greenwald et al. (1998) found race IATs to be correlated at r = .13 with feeling thermometers and r = .21 with semantic differentials. Using the same measures we found similar correlations (rs = .23 and .24, respectively).

DISCUSSION

Automatic Associations Reveal White Preference Regardless of Familiarity with Race Stimuli

The reported experiment examined whether automatic White preference emerges from greater familiarity with stimuli representing European- than African-Americans. Differential familiarity had provided a viable alternative interpretation of previous research attempting to measure attitudes using the IAT. The present findings refute the familiarity explanation by demonstrating automatic White preference even when subjective familiarity with Black and White exemplars is equated.

Besides this experiment, the influence of stimulus familiarity on automatic evaluations has also been addressed by other research providing converging evidence that stimulus familiarity does not necessarily produce automatic liking. For instance, Ottaway et al. (in press) demonstrated that Caucasian subjects expressed automatic preference for European- over African- and Hispanic-Americans even when the familiarity and frequency of all race stimuli were controlled. Likewise, Rudman, Greenwald, Mellott, and Schwartz (1999) showed that Americans exhibited greater automatic liking toward unfamiliar American than familiar Russian presidents. When automatic attitudes toward academic disciplines (e.g., mathematics) was examined, Nosek, Banaji, and Greenwald (1998) found that familiarity with math-related stimuli did not necessarily produce liking for the discipline. Specifically, they obtained predictable sex differences in attitudes toward mathematics, especially when mathematics was contrasted with unfamiliar places in the IAT. That is, women favored unfamiliar places over mathematics, whereas men favored mathematics over unfamiliar places; however, both sexes liked mathematics less than the humanities. Finally, when attitudes toward flowers and insects were examined, subjects expressed strong automatic preference for flowers over insects even when both types of stimuli were selected to be equally frequent and presumably equally familiar (Ottaway et al., in press). These data, together with our findings, suggest that the IAT can detect individual differences in automatic attitudes independent of subjective familiarity with specific stimuli.

To emphasize the extent to which the familiarity explanation has been addressed, we also briefly mention another experiment in which the familiarity explanation was challenged in a different way (Dasgupta, McGhee, Greenwald, & Banaji, 1999). In this experiment the frequency of Black and White names was

manipulated across four IATs to test whether automatic White preference emerges from frequent exposure to, and hence greater familiarity with, White stimuli. Name-frequency data were obtained from the Internet and their validity verified using name frequencies from the 1990 U.S. census. In addition to objective frequencies, subjective frequency ratings were also obtained for the same stimuli. In all, this study employed three measures of stimulus frequency (Internet-based, census-based, and frequency ratings) in order to provide a stringent test of the influence of stimulus frequency on automatic race attitudes. Contrary to the frequency explanation, a substantially stronger White preference effect was obtained when the frequency of Black and White names were controlled than when popular White and rare Black names were contrasted. By manipulating stimulus familiarity across several IATs, this study revealed yet again that stimulus familiarity does not provide a viable alternative explanation for the race IAT effect.

Converging Evidence from Pictorial and Name Stimuli

Our prediction that automatic race evaluations reveal underlying attitudes was also supported by agreement between picture and name IATs. As predicted, even when name stimuli were replaced with equally unfamiliar Black and White faces, strong pro-White attitudes were revealed. The magnitude of race preference observed across the picture and name IATs was substantial (average $d = .73$).

Although both pictures and names revealed similar attitudes, the effect was larger for names than pictures (average ds $= .93$ vs .53). This difference may have occurred for several reasons. First, faster latencies in the picture IAT suggest that pictures are easier to process than words, which reduces task difficulty. There is some indication that as the difficulty of the IAT decreases, so does the magnitude of the obtained IAT effect (Greenwald, 1999). Second, because the picture IAT required responses to both pictures and words, switching from one processing mode to another (e.g., picture to word) may have attenuated the automatic effect (cf. Park & Gabrieli, 1995; Sperber, McCauley, Ragain, & Weil, 1979). Third, our pictures depicting well-dressed attractive Black individuals may have increased the accessibility of a favorable Black subtype (e.g., young professionals; cf. Devine & Baker, 1991) and reduced the magnitude of White preference. Despite the effect size difference between name and picture IATs in the present research, it is important to underscore that both tasks revealed substantial pro-White attitudes.

Although research on beliefs and attitudes has usually depended on direct or self-report measures, techniques such as the IAT and similar methods providing indirect assessments of attitude strength may be necessary to examine socially sensitive attitudes. Ruling out the IAT's vulnerability to stimulus familiarity strengthens the argument that automatic associations captured by this technique represent underlying attitudes and can supplement traditional measures of race attitudes.

REFERENCES

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

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

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

Bellezza, F. S., Greenwald, A. G., & Banaji, M. R. (1986). Words high and low in pleasantness as rated by male and female college students. *Behavior Research Methods, Instruments, and Computers,* **18,** 299–303.

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

Crosby, F. J., Bromley, S., & Saxe, L. (1980). Recent unobtrusive studies of black and white discrimination and prejudice: A literature review. *Psychological Bulletin,* **87,** 546–563.

Dasgupta, N., McGhee, D. E., Greenwald, A. G., & Banaji, M. R. (1999). *Automatic preference for White Americans: Does the frequency of race stimuli produce the effect?* Unpublished manuscript, New School University.

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

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

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

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

Fiske, S. T. (1998). Stereotyping, prejudice, and discrimination. In D. T. Gilbert & S. T. Fiske (Eds.), *The handbook of social psychology* (Vol. 2, pp. 357–411). Boston, MA: McGraw–Hill.

Greenwald, A. G. (1999). [The magnitude of IAT effects varies as a function of task difficulty]. Unpublished raw data.

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review,* **102,** 4–27.

Greenwald, A. G., Klinger, M. R., & Schuh, E. S. (1995). Activation by marginally perceptible ("subliminal") stimuli: Dissociation of unconscious from conscious cognition. *Journal of Experimental Psychology: General,* **124,** 22–42.

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

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, IL.

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

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

Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (1998). *Gender differences in implicit attitudes toward mathematics*. Paper presented at the annual meeting for the Midwestern Psychological Association, Chicago, IL.

Ottaway, S. A., Hayden, D. C., & Oakes, M. A. (in press). Implicit attitudes and racism: The effect of word familiarity and frequency in the Implicit Association Test. *Social Cognition.*

Park, S. M., & Gabrieli, J. D. E. (1995). Perceptual and nonperceptual components of implicit

memory for pictures. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **21,** 1583–1594.

Rudman, L. A., Greenwald, A. G., Mellott, D. S., & Schwartz, J. L. K. (1999). *Implicit cognition in intergroup relations research: Construct validation for the implicit association test.* Manuscript submitted for publication.

Schuman, H., Steeh, C., & Bobo, L. (1997). *Racial attitudes in America: Trends and interpretations.* Cambridge, MA: Harvard Univ. Press.

Sperber, R. D., McCauley, C., Ragain, R. D., & Weil, C. M. (1979). Semantic priming effects on picture and word processing. *Memory and Cognition*, **7,** 339–345.

U.S. Department of Justice. (1998). [online] Database of first names from the 1990 U.S. census. Available at http://www.census.gov/genealogy/names/. [1998, January 7]

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.

Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology, Monographs*, **9** (2, Pt. 2).