Published in final edited form as:

Soc Psychol (Gott). 2010 March 1; 41(3): 152–157. doi:10.1027/1864-9335/a000022.

Explicit Control of Implicit Responses: Simple Directives can alter IAT Performance

Matthew Wallaert,
Swarthmore College

Andrew Ward, and
Swarthmore College

Traci Mann
University of Minnesota

Abstract

Research has begun to reveal the malleability of implicit prejudice. One measure of this construct, the race Implicit Association Test (IAT), represents a widely-used tool to assess individuals' positive and negative associations with different racial groups. In two studies, we demonstrate the capacity of salient pressures to alter implicit racial responses. In Study 1, an enhancement of promoting pressures through an explicit instruction to stereotype was sufficient to increase pro-White bias on the IAT. In Study 2, an enhancement of inhibiting pressures through a simple instruction to avoid stereotyping was sufficient to reduce pro-White bias. Taken together, the studies suggest that implicit prejudice is amenable to voluntary control through the use of simple, direct means.

Keywords

Implicit prejudice; IAT; attention; self-control

A growing body of literature suggests that so-called automatic stereotypes and prejudice are at least partially amenable to voluntary control (Blair, 2002; Blair, Ma, & Lenton, 2001; Devine, Plant, Amodio, Harmon-Jones, & Vance, 2000; Fazio, Jackson, Dunton, & Williams, 1995; Moskowitz, Gollwitzer, Wasel, & Schaal, 1999). Such constructs are typically assessed through the use of implicit measures, which have increased in use dramatically in the last decade (Fazio & Olson, 2003). One purported advantage of these measures over traditional explicit assessments is that they allow for the measurement of constructs that individuals may be unwilling or unable to report (Greenwald & Banaji, 1995). Perhaps the best known of this new class of measurement tools is the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998), in which individuals are required to classify four categories of stimuli using two dimensions. One form of the race IAT, for example, asks participants to classify words as good or bad and faces as belonging to Black or White racial groups. Implicit group preference is reflected in the speed with which one can, for example, respond to so-called compatible pairings (positive words and White faces and/or negative words and Black faces) faster than incompatible pairings (positive words and Black faces and/or negative words and White faces).

Automatic and Controlled Responses to the IAT

Recent theorizing has suggested that performance on the IAT in particular may represent the product of both automatic and controlled cognitive processes, in which automatic associations, to the extent they occur, can be at least partially overridden by conscious processes (Conrey, Sherman, Gawronski, Hugenberg, & Groom, 2005). Regardless of the degree of controllability of IAT responses, however, to date the evidence for the impact of experimental manipulations on the malleability of race-IAT responses has been mixed. On the one hand, manipulations involving “indirect means” have been successful in altering IAT responses (Nosek, Greenwald, & Banaji, 2006). For example, Dasgupta and Greenwald (2001) reduced pro-White bias by first exposing participants to images of respected African-American figures. Lowery, Hardin, and Sinclair (2001) achieved a similar reduction through the use of a Black rather than a White experimenter (see also Foroni & Mayr, 2005; Rudman, Ashmore, & Gary, 2001; Sinclair, Lowery, Hardin, & Colangelo, 2005). On the other hand, the results for more direct manipulations have been less positive, leading the IAT developers to conclude that “it is easier to shift IAT effects through indirect means... than through a request to deliberately alter the effects” (Nosek et al., 2006, p. 20). For example, investigations of the efficacy of “faking” responses to the IAT have found minimal supportive evidence (Banse, Seise, & Zerbes, 2001; Kim, 2003; Schnabel, Banse, & Asendorpf, 2006; Steffens, 2004), except in cases in which the mechanics of the test have been made at least somewhat apparent to participants (Fiedler & Bluemke, 2005; Kim, 2003). As Greenwald, Poehlman, Uhlman, and Banaji (2009) have stated, “Subjects who are explicitly instructed to slow their responding in one of the IAT’s two combined tasks can use that instruction to produce faked scores. At the same time, most naive subjects do not spontaneously discover this strategy” (p. 18).

Such conclusions notwithstanding, there is emerging evidence that individuals may be able to alter their responses to the IAT when explicitly instructed to do so. For example, in a study conducted by Lowery et al. (2001; Study 3), non-Black participants were told that “your job is to be as nonprejudiced as possible” (p. 848). The study showed a significant reduction in pro-White responses. However, it relied on a paper-and-pencil version of the IAT and did not employ the latest scoring algorithm (Greenwald, Nosek, & Banaji, 2003). In a similar vein, Kim (2003, p.) instructed White and Asian participants to treat the IAT “as if it may indicate that you possess prejudice but you prefer not to give that indication.” Such an instruction, though perhaps less directive than that employed by Lowery et al. (2001), also produced diminished pro-White responses (as assessed using a standard effect size measure) relative to an uninstructed control group, but the effect was not statistically significant. In addition, similar to the Lowery et al. study, the study did not rely on the latest scoring algorithm. Accordingly, the present studies were designed to explore whether simple, explicit instructions can in fact substantially reduce (or, for that matter, increase) white preference on a contemporary race IAT in the absence of any information regarding the underlying mechanics of the test (e.g., reaction times, compatible vs. incompatible blocks, etc.; cf. Fiedler & Bluemke, 2005; Kim, 2003).

Conflicting Pressures and the IAT

For many individuals, completing measures assessing implicit prejudice represents a self-control challenge (Carver & Scheier, 1981; Plant & Devine, 2009) because their responses are subject to conflicting pressures—pressures that both promote and inhibit behavior. On the one hand, many respondents are motivated to avoid the appearance of prejudice (Crandall & Eshleman, 2003; Greenwald, Nosek, & Sriram, 2006). However, the relatively automatic associations revealed by the IAT have the potential to suggest bias against a particular racial group, at least in the minds of some perceivers (but see Arkes & Tetlock, 2004). In particular, non-Black Americans’ responses to the IAT may represent a “quasi-

stationary equilibrium” between conflicting forces (Lewin, 1943; see also Mann & Ward, 2004, 2007), involving both pressures promoting the automatic expression of perceived anti-Black sentiments and inhibiting pressures representing conscious attempts to suppress such sentiments.

The notion that IAT responses reflect the result of conflicting pressures suggests at least one strategy for altering pro-White/anti-Black sentiments on the test. Namely, one could attempt to change the strength of one of the competing forces governing relevant responses. To the extent that voluntary control can influence IAT responses, one could, for example, attempt to increase the level of White preference exhibited on the test—something we asked participants to do in Study 1. In Study 2, by contrast, we investigated the capacity of a manipulation to increase the *inhibition* of pro-White/anti-Black sentiments, a result with practical implications beyond an enhanced understanding of the dynamics underlying IAT responses.

Study 1

In the first study, we investigated the capacity of an explicit instruction to alter responses on the IAT. Half the participants completed the standard race IAT. For the remaining half of participants, we attempted to increase pro-White responses by asking to them to respond to the IAT in a manner intended to increase scores reflecting White rather than Black preference. Pilot testing had revealed reluctance on the part of some participants to exhibit strong White preference on the IAT. Accordingly, measures were taken to reduce participant reactivity, as detailed below.

Method

Participants—Sixty non-Black participants completed the study in partial fulfillment of an introductory psychology course requirement.

Procedure—Participants were tested individually and told that they would be completing a reaction time task. The task involved two IATs, both run on a PC laptop computer, downloaded from the website of Dr. Brian Nosek (<http://projectimplicit.net/nosek/iat/>) and slightly altered to match the parameters detailed on the website of Dr. Anthony Greenwald (http://faculty.washington.edu/agg/iat_materials.htm) as of November 15, 2005 (including the correction to the published syntax [Greenwald et al., 2003] for addressing error penalties). The experimenter directed each participant to read the on-screen instructions and then explained that the first IAT was intended to familiarize them with the design of the experiment and would not be scored. After soliciting questions, the experimenter left the room. Following the completion of a practice IAT, involving the pairing of positive and negative words with the names of flowers and bugs, participants were instructed to summon the experimenter, who then set up the computer program to run the race IAT. Importantly, research has shown that experience with a warmup IAT does not, in and of itself, influence participants’ ability to alter performance on a subsequent IAT (Fiedler & Bluemke, 2005). One version of the race IAT required individuals to press one key when they saw either a depiction of a White face or a positive word and a second key when they saw a Black face or negative word, followed by the requirement to perform the complementary task, pairing a White face with a negative word and a Black face with a positive word. In the second version, the order of the two complementary tasks was reversed.¹

¹Relevant analyses revealed no main effects or interactions involving task order in either study.

At this point, participants were instructed to draw an envelope from a randomized assortment and follow the directions contained inside. Half of the envelopes simply instructed participants to continue by completing the subsequent IAT. The other half directed participants to try to offer racially stereotyped responses on the subsequent IAT, an instruction justified to participants as a means of investigating the limits of the test. The instructions in this condition concluded by reinforcing the directive to participants to engage in stereotyping, asking them to try to respond to the IAT like “someone who holds a strong preference for Whites over Blacks” would respond.

Results

Scores on the race IAT were analyzed according to the improved scoring algorithm published by Greenwald et al. (2003). This analysis yields a “D-score,” similar to a standard effect size measure, in which a positive score indicates faster pairing of White faces with positive words and/or faster pairing of Black faces with negative words than pairing of White faces with negative words and/or Black faces with positive words. This analysis revealed a significant increase in the D-score for participants in the condition instructed to engage in stereotyping ($M = 0.60$, $SD = 0.45$), relative to those in the control condition ($M = 0.39$, $SD = 0.35$), $t(58) = 2.02$, $p < .05$, $d = 0.52$ (see Figure 1). An inspection of the means in Table 1 indicates that this effect was associated with an increase in reaction times reflecting responses to incompatible trials, that is, slower responses on trials in which the task involved pairing of positive words with Black faces and negative words with White faces.² Of importance, this increase in reaction times in the stereotyping condition was not accompanied by a decrease in the relevant pooled standard deviations (see Greenwald et al., 2003)—a decrease that could have increased D-scores through artifactual means (see Table 1). Moreover, despite the slower responses to incompatible trials, the overall reactions times of participants instructed to attempt to enhance pro-White responses ($M = 1047$ msec, $SD = 575$ msec) did not differ significantly from those in the control condition ($M = 948$ msec, $SD = 504$ msec), $t < 1$, ns .

Discussion

An increase in the pressure to engage in stereotyping, created by an explicit instruction to offer such responses, resulted in an increase in IAT scores revealing White preference. In Study 2, we attempted to diminish White preference through an enhancement of the ability to inhibit stereotyping.

Study 2

In this study, half the participants completed the race IAT under standard conditions; the remaining half completed it immediately after being warned to avoid stereotyping. This warning took the form of a simple, straightforward admonition.

Method

Participants—Fifty-three non-Black undergraduates completed the study in partial fulfillment of an introductory psychology course requirement.

²In a further investigation of potential mechanisms underlying our results, the IAT data for Study 1 were analyzed using the Quad Model (Conrey et al., 2005), a multinomial model that uses error rates (i.e., the percentage of trials in which an incorrect pairing occurs) to estimate the rate at which implicit stereotypes are activated by IAT stimuli, as well as the rate at which activated stereotypes are overcome by deliberate, controlled responses. For the control group, participants' mean error rate was approximately 5%, and the model fit the data well, $\chi^2(1) = 0.0005$, $p = .98$. However, for the group instructed to enhance white bias, the mean error rate was 8%, and the model did not fit the data as well, $\chi^2(1) = 3.89$, $p < .05$.

Procedure—After completing the warmup “flower-bug” IAT, participants in the no-stereotype condition were told the following: “Please be careful not to stereotype on the next section of the test.” Participants in the control condition were simply told that the test would continue. The experimenter then exited the room while participants completed a race IAT identical to that described in Study 1, after which he returned and debriefed them.

Results

The mean D score for participants in the control condition was 0.42 ($SD = 0.34$), indicating a significant preference for White faces over Black faces. For those participants instructed to avoid stereotyping, the mean D score was 0.22 ($SD = 0.35$), a significant reduction in White preference, $t(51) = 2.09, p < .05$, Cohen’s $d = .58$, though still significantly above a D-score of 0, $t(23) = 3.08, p < .01$ (see Figure 1). An inspection of the means in Table 1 indicates that this effect was associated with an increase in reaction times reflecting responses to compatible trials – that is, trials in which the task involved pairing negative words with Black faces and positive words with White faces.³ In other words, when instructed to avoid stereotyping, participants increased the latency associated with negative responses to Blacks and/or positive responses to Whites. Of importance, the significant decrease in the D-score in the no stereotype condition was not accompanied by a significant increase in the relevant pooled standard deviations, $t(51) = 0.57, p = .57$ — an increase that could have decreased D-scores through artifactual means (see Table 1). In addition, the overall reactions times of participants instructed to attempt to avoid stereotyping ($M = 974$ msec, $SD = 476$ msec) did not differ significantly from those in the control condition ($M = 946$ msec, $SD = 424$ msec), $t < 1, ns$.

Discussion

A simple admonition to avoid stereotyping was sufficient to reduce pro-White responses among non-Black participants. In the past, manipulations involving “indirect means” have been successful in altering IAT responses (Nosek et al., 2006). To our knowledge, the study presented here represents the first instance of the use of a simple, direct instruction to reduce pro-White responses on a race IAT scored using the latest scoring algorithm.

Even with the benefit of an injunction, however, participants still exhibited statistically significant pro-White responses, supporting the putative dominance of relatively automatic associations. In other words, such a finding may suggest an upper bound on the capacity of individuals holding implicit White preferences to control IAT responses, assuming that individuals possessing no bias would provide an IAT D-score essentially equivalent to zero. However, such a claim has been challenged by Blanton and Jaccard (2006), who argue that it is unknown whether an individual who is “racially unbiased” would offer responses on the IAT that fall perfectly between hypothetical D-scores signifying pro-Black and pro-White bias. At this point, what is certain is that simply instructing participants to avoid stereotyping can significantly reduce the extent to which they appear to hold anti-Black sentiments as assessed by a widely used implicit measure.

General Discussion

Across two studies, the degree of White preference exhibited by non-Black participants completing the race IAT was shown to be highly malleable. Inspired in part by previous work exploring the disproportionate influence of salient pressures when completing an

³Again an effort to apply the Quad Model to our data revealed a good fit of the model to data provided by control condition participants, $\chi^2(1) = 1.07, p = .30$ (error rate = approximately 4%) but not to data provided by experimental group participants, $\chi^2(1) = 6.18, p = .01$ (error rate also = approximately 4%).

attentionally demanding task (e.g., Ward & Mann, 2000), we found that a simple directive to avoid stereotyping was sufficient to reduce scores indicating White preference, just as an instruction to engage in pro-White/anti-Black responses increased White preference. The fact that such scores could be manipulated in such a straightforward manner is somewhat surprising in light of the fact that IAT responses, especially those revealing potential prejudice, are considered by some to be beyond deliberate conscious control (see Gawronski, LeBel, & Peters, 2007). The Study 2 results suggest a capacity to consciously inhibit implicit responses can be jumpstarted by warning participants before the task has begun to be careful not to stereotype. Indeed, such pre-task instructions may signal to participants that there is even a possibility of exerting explicit control over prejudiced responses on the IAT (see Fazio & Towles-Schwen, 1999), something that participants might not otherwise discover until it is “too late.”

Further research will shed light on the mechanism underlying the effects documented in the studies reported here. Gawronski et al. (2007) have argued that efforts to understand what is assessed by implicit measures should distinguish between stereotype activation, in which relevant associations are brought to mind, and stereotype validation, in which such associations are explicitly endorsed or rejected. The fact that relevant participants in our studies showed some slowing on particular reaction time measures implicates the validation process, suggesting that they attempted to override an activated construct, either by promoting (Study 1) or inhibiting (Study 2) its expression. By contrast, the fact that these same participants did not show decreased reaction times on complementary reaction time measures is less revealing, as participants in the control condition may have already been at ceiling in terms of their ability to respond to the task as quickly as possible (see Kim, 2003). In addition, overall reaction times were not significantly altered by the manipulation employed in either study, suggesting the possibility that the activation of the relevant construct itself may have been at least somewhat altered by the instructions to increase or decrease stereotypical responses. Regardless of the particular mechanism, it is worth reemphasizing the primary conclusion from these studies—a conclusion that, to our knowledge, has not previously been demonstrated: A simple, straightforward directive, without any mention of reaction times or the underlying structure of the relevant test, can significantly alter scores on the most popular implicit measure of racial attitudes in use today.

Conclusion

Analysis of the underlying and often conflicting dynamics that shape social responses such as stereotyping and prejudice promises not only to advance our understanding of the relevant behavior but also points the way to overcoming a source of human conflict and misery. The results of our second study are particularly encouraging, for they suggest that a warning to avoid stereotyping can be an effective means for reducing expression of an implicit bias that many individuals would prefer not to possess. The simple instruction to “be careful not to stereotype” reduced white preference on the IAT by nearly 50% – a significant result, both in terms of statistical considerations and, perhaps more importantly, potential real-world application.

Acknowledgments

This research was supported by NIH Grant MH63795 to Traci Mann and Andrew Ward. We thank Frank Durgin and Jeff Sherman for their advice in analyzing the results of these studies.

References

- Arkes HR, Tetlock PE. Attributions of implicit prejudice, or “Would Jesse Jackson ‘fail’ the Implicit Association Test?” *Psychological Inquiry*. 2004; 15:257–278.
- Banise R, Seise J, Zerbes N. Implicit attitudes towards homosexuality: Reliability, validity, and controllability of the IAT. *Zeitschrift für Experimentelle Psychologie*. 2001; 48:145–160.
- Blair IV. The malleability of automatic stereotypes and prejudice. *Personality and Social Psychology Review*. 2002; 6:242–261.
- Blair IV, Ma JE, Lenton AP. Imagining stereotypes away: The moderation of implicit stereotypes through mental imagery. *Journal of Personality and Social Psychology*. 2001; 81:828–841. [PubMed: 11708560]
- Blanton H, Jaccard J. Arbitrary metrics in psychology. *American Psychologist*. 2006; 61:27–41. [PubMed: 16435974]
- Carver, CS.; Scheier, MF. *Attention and self-regulation: A control theory approach to human behavior*. New York: Springer-Verlag; 1981.
- Conrey FR, Sherman JW, Gawronski B, Hugenberg K, Groom C. Separating multiple processes in implicit social cognition: The Quad-Model of implicit task performance. *Journal of Personality and Social Psychology*. 2005; 89:469–487. [PubMed: 16287412]
- Crandall CS, Eshleman A. A justification-suppression model of the expression and experience of prejudice. *Psychological Bulletin*. 2003; 129:414–446. [PubMed: 12784937]
- Dasgupta AG, Greenwald AG. Exposure to admired group members reduces automatic intergroup bias. *Journal of Personality and Social Psychology*. 2001; 81:800–814. [PubMed: 11708558]
- Devine PG, Plant EA, Amodio DM, Harmon-Jones E, Vance S. The regulation of explicit and implicit race bias: The role of motivations to respond without prejudice. *Journal of Personality and Social Psychology*. 2002; 82:835–848. [PubMed: 12003481]
- Fazio RH, Jackson JR, Dunton BC, Williams CJ. Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*. 1995; 69:1013–1027. [PubMed: 8531054]
- Fazio RH, Olson MA. Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology*. 2003; 54:297–327.
- Fazio, RH.; Towles-Schwen, T. The MODE model of attitude-behavior processes. In: Chaiken, S.; Trope, Y., editors. *Dual process theories in social psychology*. New York: Guilford; 1999. p. 97–116.
- Fiedler K, Bluemke M. Faking the IAT: Aided and unaided response control on the implicit association tests. *Basic and Applied Social Psychology*. 2005; 27:307–316.
- Foroni F, Mayr U. The power of a story: New, automatic associations from a single reading of a short scenario. *Psychonomic Bulletin and Review*. 2005; 12:139–144. [PubMed: 15948288]
- Gawronski B, LeBel EP, Peters KR. What do implicit measures tell us?: Scrutinizing the validity of three common assumptions. *Perspectives on Psychological Science*. 2007; 2:181–193.
- Greenwald AG, Banaji MR. Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*. 1995; 102:4–27. [PubMed: 7878162]
- Greenwald AG, McGhee DE, Schwartz JKL. Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*. 1998; 74:1464–1480. [PubMed: 9654756]
- Greenwald AG, Nosek BA, Banaji MR. Understanding and using the Implicit Association Test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology*. 2003; 85:197–216. [PubMed: 12916565]
- Greenwald AG, Nosek BA, Sriram N. Consequential validity of the Implicit Association Test. *American Psychologist*. 2006; 61:56–61. [PubMed: 16435977]
- Kim DY. Voluntary controllability of the Implicit Association Test (IAT). *Social Psychology Quarterly*. 2003; 66:83–96.
- Lewin, K. *Bulletin of the National Research Council*. Vol. 108. Committee on Food Habits, National Research Council; Washington, D.C: 1943. Forces behind Food Habits and Methods of Change. The Problem of Changing Food Habits; p. 35–65.

- Lowery BS, Hardin CD, Sinclair S. Social influence effects on automatic racial prejudice. *Journal of Personality and Social Psychology*. 2001; 81:842–855. [PubMed: 11708561]
- Mann T, Ward A. To eat or not to eat: Implications of the attentional myopia model for restrained eaters. *Journal of Abnormal Psychology*. 2004; 113:90–98. [PubMed: 14992661]
- Mann T, Ward A. Attention, self-control, and health behaviors. *Current Directions in Psychological Science*. 2007; 16:280–283.
- Moskowitz GB, Gollwitzer PM, Wasel W, Schaal B. Preconscious control of stereotype activation through chronic egalitarian goals. *Journal of Personality and Social Psychology*. 1999; 77:167–184.
- Nosek, BA.; Greenwald, AG.; Banaji, MR. The Implicit Association Test at Age 7: A methodological and conceptual review. In: Bargh, JA., editor. *Automatic processes in social thinking and behavior*. Psychology Press; 2006.
- Plant EA, Devine PG. The Active Control of Prejudice: Unpacking the Intentions Guiding Control Efforts. *Journal of Personality and Social Psychology*. 2009; 96:640–652. [PubMed: 19254109]
- Rudman LA, Ashmore RD, Gary ML. Unlearning automatic biases: The malleability of implicit stereotypes and prejudice. *Journal of Personality and Social Psychology*. 2001; 81:856–868. [PubMed: 11708562]
- Schnabel K, Banse R, Asendorpf J. Employing automatic approach and avoidance tendencies for the assessment of implicit personality self-concept. *Experimental Psychology*. 2006; 53:69–76. [PubMed: 16610274]
- Sinclair S, Lowery B, Hardin CD, Colangelo A. Social tuning of automatic racial attitudes: The role of affiliative motivation. *Journal of Personality and Social Psychology*. 2005; 89:583–592. [PubMed: 16287420]
- Steffens MC. Is the Implicit Association Test immune to faking? *Experimental Psychology*. 2004; 51:165–179. [PubMed: 15267125]
- Ward A, Mann T. Don't mind if I do: Disinhibited eating under cognitive load. *Journal of Personality & Social Psychology*. 2000; 78:753–763. [PubMed: 10794378]

Figure 1.
Mean IAT “D-score” across two studies. Error bars represent standard errors.

TABLE 1**Mean Latencies By Condition**

	Mean Latency (msec)	Mean Pooled SD	Latency Difference (Incompatible – Compatible)
Study 1			
Control			
Compatible	867.77	406.30	160.19
Incompatible	1027.96		
Stereotype			
Compatible	860.50	495.09	372.07
Incompatible	1232.57		
Study 2			
Control			
Compatible	854.36	408.95	182.61
Incompatible	1036.97		
Don't Stereotype			
Compatible	920.97	445.31	105.20
Incompatible	1026.17		