The Psychological Record, 2010, 60, 57–66

THE IMPLICIT RELATIONAL ASSESSMENT PROCEDURE: EXPLORING THE IMPACT OF PRIVATE VERSUS PUBLIC CONTEXTS AND THE RESPONSE LATENCY CRITERION ON PRO-WHITE AND ANTI-BLACK STEREOTYPING AMONG WHITE IRISH INDIVIDUALS

Dermot Barnes-Holmes, Aisling Murphy, and Yvonne Barnes-Holmes
National University of Ireland, Maynooth

Ian Stewart
National University of Ireland, Galway

The current research comprised two experiments that employed the Implicit Relational Assessment Procedure (IRAP) as a measure of implicit racial attitudes. White Irish participants were exposed to blocks of trials that involved responding in a manner consistent with either a pro-white stereotype or a pro-black stereotype. In Experiment 1, participants completed the IRAP in either a public or private assessment situation. It was hypothesized that implicit pro-white stereotyping would decrease in the public context relative to the private context. The results, however, were not in accordance with this prediction. A second experiment was conducted to determine if requiring participants to respond in a public context but within a shorter timeframe would impact significantly upon implicit stereotyping. The results showed that a reduction in response latency significantly increased ingroup stereotyping. The findings appear to be consistent with the relational elaboration and coherence model.
Key words: Implicit, racism, adults, assessment context, response latency

Numerous studies using the well-known Implicit Association Test, or IAT, have indicated that white participants tend to show a relatively strong pro-white/anti-black bias (see Dasgupta, McGhee, & Greenwald, 2000; Greenwald et al., 2002; Monteith, Voils, & Ashburn-Nardo, 2001; Livingston, 2002; Ottaway, Hayden, & Oakes, 2001). In the study conducted by Dasgupta et al., for example, participants were required on some blocks of trials to categorize unfamiliar black faces or names typical of black people together

Preparation of this article was supported by funding awarded to the first, third, and fourth authors from the Irish Research Council for Humanities and Social Sciences, under the Research Development Initiative (2008–09).

The authors would like to thank two anonymous reviewers for their comments on an earlier version of the current work.

Correspondence concerning this article should be addressed to Dermot Barnes-Holmes, Department of Psychology, National University of Ireland, Maynooth, Co. Kildare, Ireland. E-mail: Dermot.Barnes-Holmes@nuim.ie

with unpleasant words, and unfamiliar white faces or names typical of white people together with pleasant words. On other blocks of trials, categorizing black stimuli with pleasant words and white stimuli with unpleasant words was required. Results showed that participants responded faster on blocks that involved categorizing pleasant words with "white" and unpleasant words with "black" than vice versa. Furthermore, this pro-white/anti-black bias occurred for participants who explicitly stated that they held no racist attitudes.

Implicit assessment methodologies such as the IAT were developed in part because explicit attitudes, often measured using self-report questionnaires, were deemed to be highly sensitive to the effects of social desirability and impression management. For example, research has shown that participants assessed in a public environment show more positive attitudes toward certain social groups on explicit measures than do participants assessed in a private environment (e.g., Blanchard, Crandall, Brigham, & Vaughn, 1994; Plant & Devine, 1998). In other words, it appears that when individuals are informed that their responses on a standard measure of racial stereotyping are going to be open to public scrutiny, rather than kept private, they tend to respond with more positive or less negative attitudes toward the outgroup. The basic assumption behind implicit measures is that they are immune, or at least far less sensitive, to such context effects. Recently, however, the IAT has been found to be susceptible to a public/private manipulation (Boysen, Vogel, & Madon, 2006; see Gawronski, LeBel, & Peters, 2007, for a review that questions the common assumption that implicit measures are immune or less sensitive to social desirability concerns).

In the study conducted by Boysen et al. (2006), the IAT was administered in both public and private assessment situations to measure bias toward homosexuality. In the public condition, participants were told that the experimenter would examine their IAT results and thus know their level of bias. In contrast, participants in the private condition were told that the experimenter would not examine their IAT performance, and thus their bias would remain unknown. The results showed that the public context significantly decreased the level of bias toward homosexuality relative to the private context. Critically, this finding suggests that the IAT may have some of the same drawbacks as explicit measures (see also Lowery, Hardin, & Sinclair, 2001; Richeson & Ambady, 2003).

An alternative implicit measure has recently been offered, the Implicit Relational Assessment Procedure (IRAP; Barnes-Holmes, Barnes-Holmes, Power, Hayden, Milne, & Stewart, 2006), which emerged from relational frame theory (RFT; Hayes, Barnes-Holmes, & Roche, 2001), a behavior-analytic account of human language and cognition. (For a brief historical review of the emergence of the IRAP from RFT, see Barnes-Holmes, Hayden, Barnes-Holmes, & Stewart, 2008; see also http://psychology.nuim.ie/IRAP/Related_Research.shtml.) The IRAP is a computer-based procedure in which participants must respond in ways that are deemed to be either consistent or inconsistent with their preexperimental learning histories. The first study to employ the IRAP (Barnes-Holmes et al., 2008) involved presenting four words on each trial: an attribute stimulus (i.e., "Pleasant" or "Unpleasant"), a positively or negatively valenced target stimulus (e.g., "Caress" or "Hate"), and two relational terms (i.e., "Similar" and "Opposite"). The response-contingent feedback for consistent blocks of trials coordinated

with previously established relations but opposed such relations during inconsistent blocks. As predicted, response latencies were shorter for consistent than for inconsistent trials (e.g., participants responded more quickly to Unpleasant-Hate–Similar than to Unpleasant–Hate-Opposite). This basic IRAP effect has now been replicated across a small number of other studies, which have shown that the IRAP (a) compares well with the IAT as a measure of individual differences (Barnes-Holmes, Murtagh, Barnes-Holmes, & Stewart, in press; Barnes-Holmes, Waldron, Barnes-Holmes, & Stewart, 2009), (b) is not easily faked (McKenna, Barnes-Holmes, Barnes-Holmes, & Stewart, 2007), (c) may be used as a measure of implicit self-esteem (Vahey, Barnes-Holmes, Barnes-Holmes, & Stewart, 2009), and (d) produces effects that clearly diverge from those obtained from explicit measures when targeting socially sensitive attitudes (Power, Barnes-Holmes, Barnes-Holmes, & Stewart, 2009).

Before proceeding, it is important to note that in creating the IRAP we sought to develop a measure that offered possible advantages over the IAT and related methods. One widely recognized limitation to the IAT is that it provides a measure of relative associative strength, which can obfuscate the exact nature of the attitudes under study (De Houwer, 2002; Nosek, Greenwald, & Banaji, 2004). If an IAT effect indicates that participants respond more quickly when pictures of white people are paired with positive stimuli and black people with negative stimuli than vice versa, this result could reflect a range of different attitudes. For example, it could indicate that white and black people are both liked, but white people are liked more than black people, or it could indicate that both white and black people are disliked but black people are disliked more than white people. In effect, the IAT can indicate that x is preferred to y, but it cannot reveal to what extent x and y are liked or disliked individually.

Alternative methods for assessing implicit cognition have been developed that aim to assess implicit attitudes toward individual concepts, such as the Go/No-Go Association Task (GNAT; Nosek and Banaji, 2001), the emotional stroop (Pratto & John, 1991), evaluative priming (Fazio, Sanbonmatsu, Powell, & Kardes, 1986), and the Extrinsic Affective Simon Task (EAST; De Houwer, 2003). Critically, however, each of these methods, along with the IAT, may be considered a relatively indirect measure of implicit attitudes or beliefs. That is, none of the methodologies requires that participants engage in a task that asks them to confirm or deny, in a direct way, the attitude or belief under investigation.

In noting the indirect nature of the IAT, De Houwer (2002) argued that

> [it] does not provide a measure of beliefs, nor was it designed to do so. It can only provide an index of associations that are assumed to be involved in certain beliefs and thus indirect evidence for the presence of certain beliefs. (pp. 117–118)

In other words, if a methodology such as the IAT indicates that pictures of black people and negative stimuli are strongly associated, for example, it is then inferred that such implicit associations underlie negative beliefs about black people. Although such an inference seems reasonable, it would also seem prudent to attempt to develop additional methodologies that aim to provide relatively direct measures of implicit cognition. The IRAP may be one such method. Unlike the IAT and other indirect measures, *each* trial

of the IRAP asks participants to confirm or deny a specific attitude or belief *directly* by responding to a relation between a label stimulus and a target item. Furthermore, the structure of the IRAP appears to permit the assessment of four separate relational responses, and thus it provides a nonrelative measure (e.g., it may be used to measure *separate* reactions to white and black people).

Although a number of studies have supported the utility of the IRAP, its sensitivity to contextual variables, such as the public/private assessment manipulation, has yet to be investigated. Experiment 1 of the current study involved administering an IRAP designed to assess implicit pro-white and anti-black stereotyping among white Irish individuals. At the time of writing, there were no published studies that had examined implicit racial bias or stereotyping in Ireland using a response latency-based measure. Participants completed the IRAP in either a public or private assessment context. Across half of the IRAP trials, participants were required to confirm that pictures of white men holding guns were safe and black men holding guns were dangerous, whereas the opposite response pattern was required across the remaining trials (i.e., white = dangerous, black = safe). Consistent with previous IRAP studies, participants were required to maintain average response latency below 3 s during the initial practice blocks. The goal of the experiment was to determine if the assessment situation would impact upon the IRAP effects in a manner similar to that observed with the IAT in the Boysen et al. (2006) study (i.e., a reduction in implicit ingroup bias in the public relative to the private assessment context). The results from Experiment 1, however, were somewhat unexpected, but post-hoc analyses also indicated that even relatively small differences in response latency on the IRAP may be an important moderating variable. Experiment 2 replicated the public context condition from Experiment 1 but required participants to maintain average response latency below 2 s rather than 3 s. The findings highlight the importance of fast responding on the IRAP and are directly relevant to a number of methodological and theoretical issues that are considered in the General Discussion section.

Experiment 1

Method

Participants

To avoid the problem of self-selection in the context of studying socially sensitive attitudes, potential participants were told simply that the experiment was concerned with "the study of memory and attitudes using a reaction-time task on a computer." Only when participants agreed to participate were they told that the study was concerned with racial stereotyping, and they were then given a minimum of 24 hr to change their minds before returning to complete the experiment. (No one chose to withdraw during this period.)

Thirty-one participants, 15 men and 16 women ranging in age from 18 to 62 years (M = 29 years), completed the experiment individually in the Department of Psychology at the National University of Ireland, Maynooth.

All participants were white Irish citizens and were randomly assigned to one of two assessment contexts—public (16 participants) and private (15 participants). No inducements were offered for participation in the study. Thirty-eight individuals commenced the experiment, but the data from 7 participants were excluded because they failed to achieve predetermined performance criteria on the IRAP (described in the following section). This level of attrition is not unusual.

Materials and Apparatus

Discrimination and Diversity scales. All participants were given three explicit self-report measures to complete. The Discrimination (DS) and Diversity (DV) scales, created by Wittenbrink, Judd, and Park (1997), counted as one explicit measure. Participants were required to indicate on 5-point scales their agreement or disagreement with a total of 14 statements, with ratings from 1 (*strongly agree*) to 5 (*strongly disagree*). The DS scale consisted of 10 statements concerning beliefs about discrimination within Irish society (e.g., "These days, reverse discrimination against Whites is as much a problem as discrimination against Blacks itself"). The DV scale consisted of four statements and targeted beliefs about the value of ethnic diversity within society (e.g., "There is a real danger that too much emphasis on cultural diversity will tear Ireland apart"). The questionnaires were scored such that 1 or 2 indicated negative racial stereotyping, 4 or 5 indicated positive racial stereotyping, and 3 indicated no stereotyping.

Modified Modern Racism Scale (MMRS). Participants also completed the MMRS, a modified version of the Modern Racism Scale (MRS; McConahay, 1986). The MMRS differed from the original in that one of the statements was removed because it was considered irrelevant in relation to an Irish society. The MMRS consisted of six statements (e.g., "Blacks should not push themselves where they're not wanted"), which targeted participants' level of bias against black people. Participants had to indicate their level of agreement or disagreement with the statements on a 5-point scale, from –2 (*strongly disagree*) to +2 (*strongly agree*). Scores were later transformed to parallel the scoring of the DS and DV scales (i.e., 1 and 2, negative racial stereotyping; 3, no stereotyping; 4 and 5, positive stereotyping).

Likert scales. Participants were required to complete 13-point Likert scales indicating how safe or dangerous they felt each IRAP target stimulus looked (i.e., three pictures of black men and three pictures of white men holding guns). Each scale ranged from –6 (*extremely safe*) to +6 (*extremely dangerous*), and a total of six scales were presented, one for each IRAP target stimulus. To use the Likert scales as an explicit measure of racial stereotyping, the average ratings given for the black targets were subtracted from the average ratings given to the white targets for each participant. Thus a positive score indicated pro-black stereotyping (white more dangerous than black), and a negative score indicated pro-white stereotyping (black more dangerous than white).

Implicit Relational Assessment Procedure (IRAP). All participants completed the IRAP on a personal computer (Dell Pentium 4®). The IRAP software was used to present the stimuli and record participants' responses. Each IRAP trial presented one of two category labels, "Safe" or "Dangerous," one of six target stimuli, and two response options, "True" and "False" (e.g.,

see Figure 1). The six target stimuli were digital color photographs that had been used in a virtual-reality study of racial prejudice by Greenwald, Oakes, and Hoffman (2003). The pictures were 305 pixels tall × 259 pixels wide and in 256-color format. Three of the pictures showed black men holding guns and three showed white men holding the same guns. All six men were wearing plain white t-shirts and were standing in front of the same red-brick background. The program also presented the IRAP instructions and a consent form.

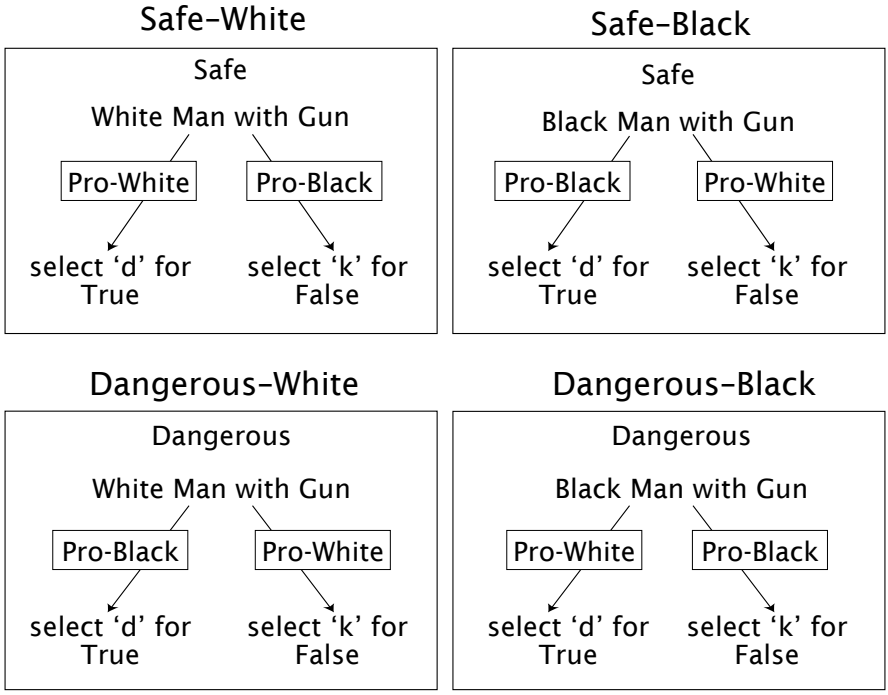

Figure 1. The four IRAP trial types. The category label ("Safe" or "Dangerous"), target stimulus (a picture of a white or black man holding a gun), and response options ("True" and "False") appeared simultaneously on each trial. Arrows with superimposed text boxes indicate which responses were deemed pro-white or pro-black. (Boxes and arrows did not appear on screen.)

Procedure

Public and private contexts. Participants in the public context were given a form consisting of a "public" statement, which they had to read and summarize:

> You are about to take a measure of racial prejudice on a computer. When you finish the test the computer will calculate the level of bias you have toward black people on a scale from 0, meaning low bias, to 100, meaning the most bias possible. After I record your computer score, your bias will also be evaluated using some surveys.

This statement was used as a tool to elicit feelings of social desirability within the public-context group, such that these participants may attempt to appear less racially biased on both the explicit measures (MMRS, DS and DV scales, Likert scales) and the implicit measure (IRAP). The participants in the private group were not given this form to fill out and were told that the experimenter would collect but not examine their scores, with the implication that individual levels of racial stereotyping would remain unknown.

The public/private manipulation was also operationalized in the following way. In the public context, the experimenter sat adjacent to the participant and watched as he or she responded to the IRAP trials.[1] In the private context, however, the experimenter left the room while the participant completed the IRAP and did not return until the computer task was finished. For the explicit measures, participants in the private context were given the three scales in booklet form and told to fill them out by circling the numbers that corresponded to their own feelings; they were also told not to mark the booklets in any other way (such as writing their names on them) because their answers were confidential. Public-context participants were given the booklet to read but had to call out their answers to the experimenter, who then wrote them down.

Implicit measure. The IRAP program began with a set of instructions, which described the task by illustrating the layout of the screen and explaining the response options (available from the first author upon request). The instructions informed participants that on each trial one of two labels, "Safe" or "Dangerous," would appear at the top of the screen along with a picture presented in the center of screen. Participants were also told that the response options "True" and "False" would appear at the bottom of the screen, and they were required to choose one of these options on each trial by pressing either the "d" or "k" key. They were told that the left–right positions of these response options would switch randomly from trial to trial. The instructions also explained that the IRAP consisted of four different trial types, and illustrated examples of these were provided. In explaining these trial types, the experimenter informed participants that sometimes they would be required to respond in a way that was consistent with their beliefs and at other times they would have to respond in a way that was inconsistent with their beliefs. Participants were assured that this was part of the experiment, and it was important for them to respond as quickly and accurately as possible on all trials of the IRAP. (At no point were participants informed which part of the experiment would be contradictory to their beliefs.) The instructions also informed participants that correct responses would allow them to progress to the next trial, but incorrect responses would produce a red X in the middle of the screen, which could only be removed by pressing the correct key.

The IRAP task consisted of a minimum of two practice blocks and a fixed set of six test blocks. Each block presented the same 24 trials,

1 In the study by Boysen et al. (2006), the experimenter did *not* sit beside the participant during the IAT. However, the IRAP typically requires more time and "cognitive effort" to complete than the IAT (approximately 10–20 vs. 5–10 min, respectively), and during pilot work for the current study some participants reported that they had "forgotten about" the public instruction by the time they had reached the test blocks. Consequently, we sought to maintain the salience of the public context throughout the IRAP with the presence of an observer.

comprised of what are defined as four different trial types (see Figure 1). The first block of the IRAP was designed to be consistent with pro-white stereotyping (e.g., Safe–White–True; Safe–Black–False; Dangerous–White–False; Dangerous–Black–True). The feedback contingencies alternated from block to block between pro-white and pro-black. Thus, in the second block, for example, the correct responses were as follows: Safe–White–False; Safe–Black–True; Dangerous–White–True; Dangerous–Black–False. Before each new block began, the participants were informed that the previously correct and incorrect answers would be reversed. The order in which IRAP blocks were presented was not counterbalanced across participants because previous research has found that this variable does not interact significantly with the critical IRAP effect (e.g., McKenna et al., 2007; Power et al., 2009; Vahey et al., 2009).

Each IRAP block consisted of 24 trials, with each target stimulus presented twice in the presence of each of the two labels. The trials were presented quasirandomly with the constraint that none of the four trial types could be presented twice in succession. The positioning of the two response options was also quasirandom in that they could not appear in the same position three times in succession.

For the first two practice blocks, participants were informed that it was a practice phase and errors were expected. Participants were required to reach a standard of ≥ 80% correct responses and a median response time of ≤ 3,000 ms. These criteria were used to ensure that participants understood and were complying with the IRAP instructions. If participants failed to achieve the two criteria for either of the two practice blocks, the required standard and the standard of responding they had achieved were presented on the screen. Participants were allowed three attempts (a total of six practice blocks) to achieve the practice criteria, and if they failed to do so, they were thanked and debriefed and their data were discarded. (Three participants were removed from the study on this basis.) Participants who did achieve the practice criteria proceeded to the six test blocks.

The procedure for the first test block was similar to the first practice block (pro-white), except that on-screen instructions informed participants that the next phase was a test and to "go quickly," although making "a few errors is okay." The second test block was similar to the second practice block (pro-black) but with the modified instructions to go quickly. Test Blocks 3 and 5 were the same as Block 1 (pro-white), and Test Blocks 4 and 6 were the same as Block 2 (pro-black). No performance criteria were applied during the test blocks in order to proceed, but if a participant's performance fell below 80% accuracy for any test block, the data for that participant were discarded. (Four participants were removed from the study on this basis.) When all six test blocks had been completed, participants reported to the researcher.

Explicit measures. Participants were given the three explicit measures to complete, the DS and DV scales, the MMRS, and the Likert scales. As noted earlier, public participants called out their answers to the experimenter, whereas the private group completed the scales unobserved by circling the appropriate numbers on the questionnaires. The participants were then thanked and debriefed, and any questions were answered. All participants completed the experiment in a single session that lasted approximately 30–40 min.

Results and Discussion

Implicit Measure

Data preparation. The primary datum was response latency, defined as the time in milliseconds that elapsed between the onset of a trial and a correct response emitted by a participant. To control for individual variations in speed of responding that may act as a possible confound when analyzing between-group differences, the response latency data for each participant were transformed into D_{IRAP} scores (Barnes-Holmes, Murtagh, et al., in press; Barnes-Holmes et al., 2009; Cullen & Barnes-Holmes, 2008; Vahey et al., 2009) using an adaptation of the Greenwald, Nosek, and Banaji (2003) *D* algorithm.

The steps involved in calculating the D_{IRAP} scores were as follows: (1) Only response latency data from the six test blocks were used; (2) latencies above 10,000 ms were removed from the data set; (3) if the data from a participant contained more than 10% of test-block trials with latencies less than 300 ms, that participant was removed from the analyses; (4) 12 standard deviations for the four trial types were calculated: 4 for the response latencies from Test Blocks 1 and 2, 4 from the latencies from Test Blocks 3 and 4, and a further 4 from Test Blocks 5 and 6; (5) 24 mean latencies were then calculated for the four trial types in each test block; (6) difference scores for each of the four trial types were calculated for each pair of test blocks by subtracting the mean latency of the pro-white test block from the mean latency of the corresponding pro-black test block; (7) each difference score was then divided by its corresponding standard deviation from step 4, yielding 12 D_{IRAP} scores, one score for each trial type for each pair of test blocks; (8) four overall trial-type D_{IRAP} scores were then calculated by averaging the scores for each trial type across the three pairs of test blocks; and (9) an overall D_{IRAP} score was calculated by averaging all 12 trial-type D_{IRAP} scores from step 7.

Data analyses. The D_{IRAP} scores for the four trial types under the public and private contexts are presented in Figure 2. The data show that Safe–White, Dangerous–White, and Dangerous–Black were in the predicted direction in both contexts, although they appeared relatively weak for the latter two trial types. That is, in general, both groups responded "True" more quickly than "False" on Safe–White and Dangerous–Black trials and responded "False" more quickly than "True" on Dangerous–White trials. Critically, the effect for the Safe–Black trial type was in the nonpredicted direction in both contexts, and paradoxically it showed a stronger effect in the private context. That is, participants responded "True" more quickly than "False" on Safe–Black trials, and the private context increased this effect.

A 2 × 4 mixed repeated-measures analysis of variance (ANOVA) was conducted on the D_{IRAP} scores, with private and public contexts as the between-participant variable and trial type as the within-participant variable. The analysis revealed a significant main effect for trial type, $F(3, 29) = 7.905$, $p < .001$, $\eta_p^2 = .13$, but no effect for context or interaction (ps > .2). Four planned-comparisons one-way between-groups ANOVAs were conducted to determine the impact of context on each of the four trial types.

The analysis for Safe–Black approached significance ($p = .08$), but the others were statistically similar (ps > .7). Eight one-sample t tests were conducted to determine if the D_{IRAP} trial-type scores differed significantly from zero. In the public context, only the Safe–White effect differed significantly ($t = 3.16$, $df = 15$, $p < .01$), and this was in the predicted pro-white direction (remaining ps > .3). In the private context, only the Safe–Black effect was significant ($t = -5.28$, $df = 14$, $p < .0001$), but it was in the opposite direction (pro-black) to that predicted (remaining ps > .13). Overall, therefore, the ANOVA for the Safe–Black trial type, combined with the results of the one-sample t tests, suggest that the private rather than the public context reduced the predicted pro-white stereotyping.

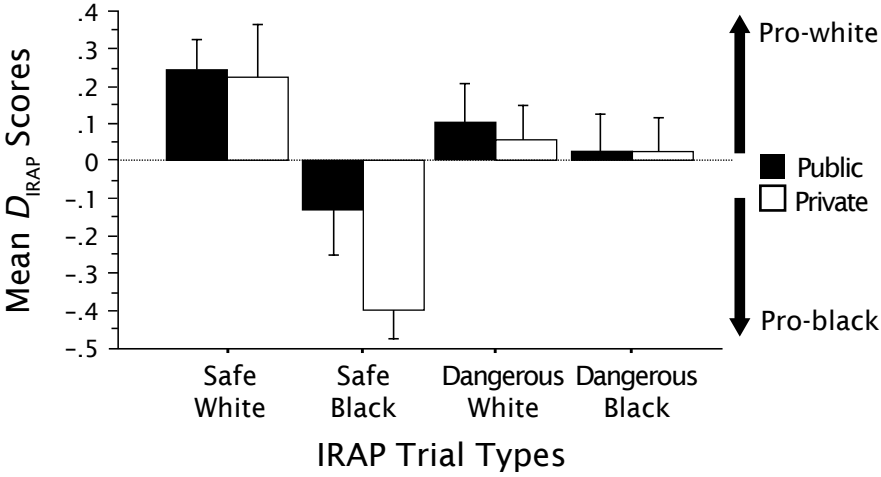

Figure 2. Mean D_{IRAP} scores for the four IRAP trial types in public and private contexts. A positive D_{IRAP} score indicates pro-white stereotyping and a negative score indicates pro-black stereotyping.

Internal reliability. In order to calculate split-half reliability for the IRAP, two overall D scores were calculated, one for odd trials and the second for even trials. These two scores were calculated in the same way as for the overall D_{IRAP} score (see step 9 above), except that the algorithm was applied separately to all odd trials and to all even trials. The split-half correlations, applying Spearman-Brown corrections, between odd and even D_{IRAP} scores, calculated separately for the two groups, proved to be moderate and significant for the public context, $r = .44$, $n = 16$, $p < .05$ (Cronbach's alpha = .45), but weak and nonsignificant for the private context, $r = .23$, $n = 15$, $p = .14$ (Cronbach's alpha = .24).

Explicit Measures

The overall mean scores on the MMRS for the private ($M = 3.9$, $SD = .6$) and public ($M = 3.9$, $SD = .7$) contexts showed that both groups exhibited positive black stereotyping. A one-way between-groups ANOVA proved to be nonsignificant ($p = .91$). The overall means for the DS and DV scales showed that participants in the public context revealed more positive racial

stereotyping across both scales (M = 3.5, SD = .6) than did participants in the private context (M = 3.1, SD = .55). A one-way between-groups ANOVA indicated that the effect was significant, $F(1, 29) = 4.6$, $p < .05$, $\eta^2 = .13$. The overall mean scores on the Likert scales showed that the public group exhibited stronger pro-black stereotyping (M = .42, SD = 2.1) than did the private group (M = .18, SD = 1.1). However, a one-way between-groups ANOVA indicated that the difference was nonsignificant (p = .6).

Implicit–Explicit Correlations

Two correlation matrices of the implicit and explicit measures were calculated: one for the public and one for the private context. Each matrix thus involved correlating the four trial-type and overall D_{IRAP} scores with each of the three explicit measures. The public context yielded a significant correlation between the Safe–White trial type and the DS and DV scales (r = –.62, n = 16, $p < .01$). In effect, larger Safe–White D_{IRAP} scores predicted increased racial bias on the explicit measure. (The correlation is negative because lower scores on the DS and DV scales indicate higher racial bias.) A similarly negative correlation between Safe–White and the MMRS scale approached significance (r = –.44, n = 16, $p < .09$), again suggesting that higher D_{IRAP} scores on this trial type predicted greater racial bias on the explicit measure. All other correlations between the five D_{IRAP} scores and the three explicit measures were nonsignificant (remaining ps > .13). The correlation matrix for the private-context group revealed no significant correlations (all ps > .16).

Summary

Contrary to experimental predictions, the public/private context manipulation had no overall significant impact on IRAP performance. In fact, paradoxically, follow-up analyses indicated that participants in the private context were found to be significantly pro-black (on the Safe–Black trial type), whereas the participants in the public context were found to be significantly pro-white (on the Safe–White trial type). Based on previous research, it was hypothesized that the public group would show less racial stereotyping than the private group, but the opposite was found to be the case. Two of the three explicit measures, however, yielded results that were broadly consistent with experimental predictions. In general, the IRAP failed to correlate with the explicit measures, although the Safe–White trial type correlated significantly with the DS and DV scales and approached significance with the MMRS, but only in the public context.

Post-Hoc Analyses

Given that the results from Experiment 1 were contrary to previously published research, it seemed wise to review the data for possible artifacts that might account for the unexpected findings. One possibility that appeared worthy of further scrutiny was the extent to which the private/public manipulation impacted upon the length of the response latencies across both consistent and inconsistent trials. Anecdotal evidence suggested that some participants may have failed to respond relatively quickly in the private context because they were not being monitored by the experimenter.

To test this post-hoc explanation, the overall mean response latencies (across all test trials) for each participant were calculated (group means: public, $M = 1{,}967$, $SD = 307$; private, $M = 2{,}462$, $SD = 413$) and were entered into a one-way between-groups ANOVA. Interestingly, a significant effect was found, $F(1, 29) = 14.5$, $p < .001$, $\eta^2 = .3$, indicating that the private group did indeed take significantly longer to respond on the IRAP trials. Given this difference, it was possible that the IRAP effects observed in the private context arose, at least in part, from reduced "automaticity" (Moors & De Houwer, 2006). In other words, the longer a participant takes to respond on the IRAP, the more unreliable it may become as a measure of implicit attitudes. (We will consider this issue in more detail in the General Discussion section.)

Experiment 2

At this point, we decided to explore the effect of time pressure on the current IRAP to determine if it does indeed have a significant impact on performance. Thus a second experiment was conducted in which participants were asked to complete the same IRAP as in Experiment 1 in a public setting, but using a 2,000-ms latency criterion for the practice blocks. These data could then be compared with the public-setting data from Experiment 1 to determine if decreasing the response latency criterion significantly impacts upon the IRAP effects. As an aside, initially we planned to include both private and public setting conditions in Experiment 2. However, pilot work showed that participants in a private setting tended to increase mean response latencies above 2,000 ms during the test blocks (presumably because performance was not being monitored), and thus a meaningful comparison with the public condition, in which mean latencies less than 2,000 ms were maintained, would not be possible. (See the General Discussion section for a recent development in the IRAP software that may help to prevent upward drift in response latencies during test blocks.)

Method

Participants

The selection technique employed in Experiment 1 was employed in Experiment 2, and again no one chose to withdraw during the 24-hr "change-of-mind" period. Nineteen white Irish citizens completed the study, 10 women and 9 men, with an age range of 18 to 52 years ($M = 27$ years). All participants completed the study in a public context. Twenty-four individuals commenced the experiment, but the data from 5 participants were excluded because they failed to achieve the predetermined performance criteria on the IRAP.

Materials and Apparatus

The apparatus and materials used in Experiment 2 were identical to those used in the public context of Experiment 1 (i.e., public statement, IRAP, MMRS, DV and DS scales, and Likert scales). Note, however, that the IRAP instructions were modified slightly to indicate that participants had to reach an average response latency of $\leq 2{,}000$ ms across each practice block of the IRAP before they could proceed to the test blocks.

Procedure

The procedure in Experiment 2 was identical to that used for the public context in Experiment 1, except that the response latency practice criterion was reduced from 3,000 to 2,000 ms, and all instructions and feedback were adjusted to reflect this change. Three participants failed to reach the practice criteria (i.e., ≥ 80% correct and a median response latency ≤ 2,000 ms) and thus did not proceed to the test blocks. The data for 2 participants were removed because their accuracy levels on one or more test blocks fell below 80% correct.

Results and Discussion

Implicit Measure

Data analyses. The response latencies were subjected to the same data preparation procedures as were employed for Experiment 1.[2] The four overall mean D_{IRAP} scores are presented in Figure 3 (2 s) and are presented, for the purposes of comparison, with the scores obtained from the public condition in the previous experiment (3 s). The results indicate that the D_{IRAP} scores for the Safe–White and Dangerous–Black trial types were considerably larger in the 2-s relative to the 3-s condition; the D_{IRAP} scores for the Safe–Black and Dangerous–White trial types were similar across conditions.

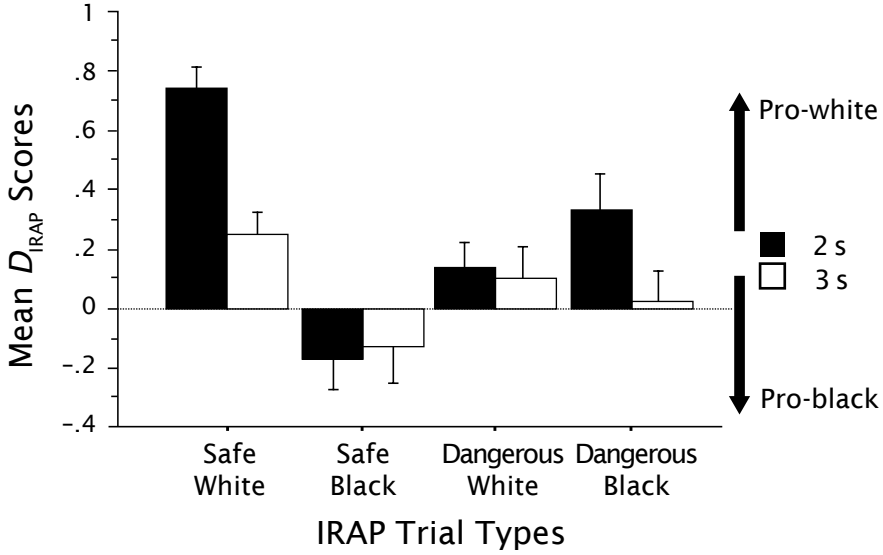

Figure 3. Mean D_{IRAP} scores for the four trial types in the 2-s and 3-s response latency conditions (public context).

2 To determine if the 2,000-ms practice criterion had impacted performance on the subsequent test blocks, the overall mean response latency calculated across all test trials was obtained, $M = 1{,}560$, $SD = 65.3$, and was compared to the overall latency obtained in the 3,000-ms public condition from Experiment 1, $M = 1{,}967$, $SD = 307$. The difference between the two practice criteria proved to be significant, $F(1, 33) = 16.5$, $p < .0004$, $\eta^2 = .3$, thus indicating that participants did in fact respond significantly more quickly during test blocks following exposure to a 2,000-ms relative to a 3,000-ms practice criterion.

A 2 × 4 mixed repeated-measures ANOVA was conducted on the D_{IRAP} scores, with 2- and 3-s conditions as the between-participant variable and trial type as the within-participant variable. There were significant main effects for time, $F(1, 33) = 6.2$, $p < .03$, $\eta_p^2 = .08$, and for trial type, $F(3, 33) = 15.35$, $p < .001$, $\eta_p^2 = .32$, and a significant interaction, $F(3, 33) = 3.4$, $p < .03$, $\eta_p^2 = .09$. Four planned-comparisons one-way between-groups ANOVAs were conducted to investigate the impact of time on each of the four trial types. Time was found to impact significantly on Safe–White, $F(1, 33) = 16.71$, $p < .001$, $\eta^2 = .34$, and Dangerous–Black, $F(1, 33) = 4.29$, $p < .05$, $\eta^2 = .12$, but not on Safe–Black or Dangerous–White trial types (ps > .7).

Four one-sample t tests were conducted to determine if the D_{IRAP} trial-type scores differed significantly from zero in the 2-s condition. (As noted previously, only the Safe–White trial type produced a significant IRAP effect in the public 3-s condition.) In the 2-s condition both Safe–White and Dangerous–Black differed significantly from zero in the predicted directions (ps < .008), and Dangerous–White approached significance ($p < .1$). In effect, decreasing the response latency practice criterion increased significantly the ingroup pro-white stereotyping (on the Safe–White trial type) and served to produce significant anti-black stereotyping (on the Dangerous–Black trial type) not observed under the 3-s latency criterion.

Internal reliability. The split-half correlation, with Spearman-Brown correction, between odd and even D_{IRAP} scores for the 2-s condition was strong and significant ($r = .81$, $p < .001$; Cronbach's alpha = .81) and was almost twice that of the public 3-s condition (.44); the difference in correlations proved to be significant ($p < .01$).

Explicit Measures

The overall mean scores from the three explicit measures for the 2-s condition were as follows: MMRS = 4.0 (*SD* = .5), DS and DV = 3.6 (*SD* = .3), and Likert = –.09 (*SD* = 1.0). Three one-way between-groups ANOVAs indicated that these scores did not differ significantly from the 3-s condition (all ps > .6), thus confirming that the IRAP response latency criterion had no impact on the explicit measures.

Implicit–Explicit Correlations

A correlation matrix of the implicit and explicit measures for the 2-s condition yielded one significant correlation between the overall D_{IRAP} score and the MMRS ($r = –.5$, $p < .03$) and a correlation that approached significance between the Safe–Black IRAP trial type and the same explicit measure ($r = –.43$, $p < .07$); all remaining ps > .13. In effect, the IRAP predicted racial bias on the MMRS, with a tentative suggestion that responding to black men holding guns as safe predicted reduced racial bias on the explicit measure. (The correlations are negative because lower scores on the IRAP and higher scores on the MMRS indicate less racial bias.)

Summary

Both the Safe–White and Dangerous–Black trial types produced IRAP effects that were significantly more pro-white in the 2-s than in the 3-s condition. Furthermore, only in the 2-s condition was anti-black stereotyping

observed (i.e., participants responded "True" more quickly than "False" on Dangerous–Black trials). The results from Experiment 2, when compared with the public condition in Experiment 1, thus confirmed the prediction that increasing automaticity, by reducing response latency, would produce greater evidence of racial stereotyping on the IRAP. (There was no obvious impact on the explicit measures.) Furthermore, the internal reliability of the IRAP in the 2-s condition was almost twice that of the 3-s condition.

General Discussion

To our knowledge, the current study is the first to examine implicit racial stereotyping against black people among white Irish individuals. The current research is also the first to explore the effect of a public versus private setting on the IRAP, and the impact of reducing response latency. Consistent with previous IAT studies investigating contextual effects (e.g., Boysen et al., 2006; Lowery et al., 2001), we hypothesized that the public group would exhibit less racial stereotyping than would the private group. However, the public versus private manipulation in Experiment 1 had no significant impact on overall IRAP performance, and an effect opposite to that expected was found at the trial-type level. That is, the private group produced significant pro-black stereotyping on the Safe–Black trial type, and the public group produced significant pro-white stereotyping on the Safe–White trial type. The unexpected results led to a post-hoc analysis of the raw IRAP data, and context was found to impact significantly upon response latency, with participants in the private context responding more slowly during the test blocks than participants in the public context. This finding suggested that participants produced less evidence of racial stereotyping in the private condition because they responded more slowly than did participants in the public condition.

The results from Experiment 2 supported the prediction that decreasing the response latency practice criterion would impact upon the overall IRAP effect. That is, significantly greater pro-white and anti-black implicit stereotyping was revealed using a 2-s rather than 3-s criterion, and, in fact, anti-black stereotyping was observed only in the 2-s condition. Furthermore, the IRAP effects obtained in Experiment 2 appear to offer increased reliability and validity because (a) internal reliability almost doubled in the 2-s condition, and (b) the results were generally consistent with previous research on implicit attitudes, which has shown that white participants tend to exhibit a pro-white/anti-black bias (e.g., Greenwald et al., 2002; Monteith et al., 2001; Ottoway et al., 2001).

The current findings are important for at least two reasons. First, they serve to highlight the necessity of relatively "fast" responding on the IRAP, at least when one is attempting to measure socially sensitive attitudes. Indeed, since collecting these data, our research group has started to employ practice criteria of 2-s or less, and the IRAP software has recently been modified to provide trial-by-trial feedback on response latency. (The message "Too Slow" now appears onscreen if a participant exceeds the response latency criterion during the practice or test blocks.) Preliminary research has been consistent with the current data in that these modifications appear to increase both IRAP effect sizes and internal reliability. It seems important, therefore, to disseminate the current findings so that other researchers employing the

IRAP are made aware of the importance of establishing and maintaining the rapid-response requirement. Indeed, given the current findings, any future IRAP study that attempts to examine the effects of a public versus private manipulation, or any other contextual variables, should insure that response latency is kept relatively constant across contexts.[3]

The second reason the current findings seem important is that relatively strong pro-white/anti-black stereotyping was observed in Experiment 2. Literature searches using the database PsychInfo® failed to yield any published study that has examined implicit racial bias or stereotyping against black people in Ireland using a response latency-based measure. Exploring the implicit attitudes of white Irish people toward black people may be of particular interest at present, because the number of black African immigrants living in Ireland has seen a sharp rise in recent years. For example, past censuses show that the number of African nationals living in Ireland increased almost tenfold from 4,867 in 1996 to 42,764 in 2006 (http://www.cso.ie/census/default.htm). Of course, racially biased responding on an IRAP may not translate into overt racist action, but an overreliance on self-report methods as a means of measuring racial bias in Ireland would seem unwise. For example, in other countries self-report measures have suggested that racist attitudes are declining, but on balance there is evidence that racial discrimination is widespread, particularly in employment situations (Brief, Dietz, Cohen, Pugh, & Vaslow, 2000; Maass, Castelli, & Arcuri, 2000). Perhaps, therefore, racial stereotyping has not been reduced to the extent implied by explicit measures alone. In any case, the current study appears to be the first to show clear evidence of implicit pro-white/anti-black stereotyping in Ireland, and thus these findings should inform any attempt to study, and hopefully undermine, racial discrimination in the Irish context.

As stated earlier, a previous study by Boysen et al. (2006) indicated that implicit bias measured with an IAT was reduced in a public relative to a private context. Unexpectedly, evidence for the opposite effect was observed in Experiment 1 in the current study, but differences in response latency appeared to play a critical role. Unfortunately, overall raw latencies were not reported by Boysen et al., and thus we cannot determine if this variable did in fact differ across contexts in the earlier study. If the latencies did differ, however, presumably they did so in a manner opposite to that observed in the current research. In making this argument, it is important to note that, unlike the IRAP, the IAT does *not* impose a response latency criterion during initial practice blocks (a simple "go fast" instruction appears before the test blocks), and thus accuracy may be a more salient feature with the IAT. Consequently, participants in the Boysen et al. study may have felt increased social pressure in the public context to maintain high levels of accuracy at the expense of slightly longer latencies, which resulted in reduced automaticity and thus less implicit bias relative to the private condition. Admittedly, the foregoing analysis is highly speculative, but given the current findings, greater attention to the possible impact of contextual

3 Although we have started to use a 2-s latency criterion, this may be too short when terms or statements, rather than single words or pictures, are being employed with the IRAP (e.g., Power et al., 2009). Furthermore, if participants are particularly young or old, unfamiliar with computers, or "challenged" intellectually in some way, we would recommend careful pilot testing to determine the minimal latency that may be employed without incurring high attrition rates.

manipulations on overall response latencies seems warranted in future IAT research.[4]

The results of the implicit–explicit correlations in the current study are generally consistent with broadly similar IAT research, in that the correlations were either weak or absent (e.g., Dasgupta et al., 2000), although the overall D_{IRAP} score correlated moderately with the MMRS in Experiment 2. The important finding in the present study, however, is that reduced response latencies increased racial stereotyping on the IRAP, while the self-report measures remained relatively unchanged. Thus, in Experiment 2, participants responded to white men holding guns as relatively safe and black men holding guns as relatively dangerous, but produced either positive or neutral racial bias on the explicit measures. How might we explain this result?

The Relational Elaboration and Coherence Model

The finding that a reduction in response latency appears to increase implicit stereotyping is consistent with a behavior-analytic and RFT account that has been offered for the IRAP effect (Barnes-Holmes et al., 2006), which we now refer to as the relational elaboration and coherence (REC) model (Barnes-Holmes, Barnes-Holmes, Stewart, & Boles, in press). The REC model assumes that specific IRAP trials may produce an immediate and relatively brief relational response before the participant actually presses a response key. The probability of this initial response will often be determined by the verbal and nonverbal histories of the participant and current contextual variables. By definition, the most probable immediate response will be emitted first most often, and thus any IRAP trial that requires a key press that coordinates with that immediate response will be emitted relatively quickly; however, if an IRAP trial requires a key press that opposes the immediate relational response, it may be emitted less quickly. Thus, across multiple trials, the average latency for inconsistent blocks will be longer than for consistent trials. In short, the IRAP effect is based on immediate relational responding, which is made apparent to the researcher when the behavioral system is put under pressure to respond quickly and accurately.

Given that pressure to respond quickly was greatest in Experiment 2 of the current study, the results indicate that the immediate relational responses Safe–White–True and Dangerous–Black–True predominated. (We shall consider the other two trial types subsequently.) According to the REC model, such response patterns would likely emerge from exposure to some of the verbal and nonverbal contingencies that operate for white individuals who are raised and live in Ireland (e.g., portrayal of young black males in the British and North American media as violent gang members). In attempting to explain why such contingencies had little if any impact on self-reports, the REC model assumes that responses to these measures likely reflected relatively elaborate and coherent relational responding. In other words, when asked to express an attitude or belief on a particular

4 An early IAT study employed a response-latency window (Cunningham, Preacher, & Banaji, 2001), but the window was relatively brief and narrow (225 to 675 ms), response-error feedback was not presented, and the researchers reported response accuracy, not latency, as the dependent variable. It is therefore difficult to make any useful comparisons between this work and the current study.

issue, it is likely that a person will produce a relational response that coheres with one or more other relational responses in his or her behavioral repertoire (see Barnes-Holmes, Hayes, & Dymond, 2001). Imagine, for example, that a participant produced the same ratings for black and white men holding guns on the semantic differentials. Such relational responses would likely cohere with *other* relevant relational networks, such as "Apart from race, the pictures are very similar" and "It is wrong to discriminate on the basis of race." The critical point here is that explicit measures are typically not completed under high time pressure, and thus participants have sufficient time to engage in the extended relational responding that is needed to produce a response that coheres with one or more other relational responses. When exposed to a time-pressured IRAP, however, the impact of a participant's elaborated relational responding would be absent or much reduced because there is insufficient time, on a trial-by-trial basis, to engage in the additional and sometimes complex relational activity that serves to generate a relationally coherent response.

In summary, therefore, the REC model assumes that the IRAP effect, when produced under appropriate time pressure, is driven largely by immediate and relatively brief relational responses, whereas explicit measures reflect extended and coherent relational networks. Or more informally, the IRAP captures spontaneous and automatic evaluations, whereas explicit measures capture more carefully considered reactions. The core of the REC model explanation for the impact of increased time pressure on the divergence between implicit and explicit measures rests on the following two assumptions. First, immediate or automatic evaluative responses may or may not cohere with subsequent relational responding. When they cohere, implicit and explicit measures will typically converge, but when they do not, the measures will typically diverge.[5] In other words, it is assumed that participants usually "reject" their immediate and brief relational responses (or automatic evaluations) if they do not cohere with their more elaborate and extended relational responding.[6] Second, the REC model predicts that the divergence between implicit and explicit "socially sensitive" attitudes should increase with greater time pressure on the IRAP, because participants have less time to engage in elaborated relational responding. In effect, as time pressure increases on the IRAP, the "contaminating" effects of elaborated relational responding on response

5 The term *diverge* is used here to indicate effects that do not go in the same direction (e.g., if a negative racial bias is observed on the IRAP but not on an explicit measure). Note, however, that when measures diverge in this way they may not correlate negatively and may even correlate positively. For example, individuals who produce high levels of negative racial bias on an IRAP may produce low levels of positive racial bias on an explicit measure, whereas individuals who produce low levels of negative implicit bias may produce high levels of positive explicit bias. Such a pattern would produce overall effects that diverge in direction on a graph but correlate positively.

6 It should be noted that the REC model does *not* predict that additional relational activity will always produce a positive response in a socially sensitive area. For some individuals, additional responding may produce a negative response that coheres with the initial negative evaluation (e.g., "The black man in the photograph looks dangerous *and* it is okay to discriminate on the basis of race"). Alternatively, additional responding may produce a relational response that allows two initially incoherent networks to cohere (e.g., "The black man looks dangerous, *but* it is wrong to discriminate on the basis of race. *However,* the black man in this *particular* photograph does look quite dangerous").

latencies decrease. Note, however, that the REC model does not predict that decreasing time pressure on the IRAP will necessarily produce increasing convergence with explicit measures. As time pressure decreases, it is difficult to predict exactly what variables will impact upon response latency, and thus the potential utility of the measure is lost. Indeed, the current findings support this conclusion because the internal reliability of the IRAP decreased as latencies increased.

As noted previously, the results from Experiment 2 produced strong IRAP effects for the Safe–White and Dangerous–Black trial types. It seems logical, therefore, to predict relatively strong effects for Dangerous–White and Safe–Black trials, but these effects were not clearly evident in the data. To explain this apparent anomaly, it is important to bear in mind that the REC model predicts IRAP effects based on immediate and brief relational responses, not logical reasoning, which typically involves relatively elaborate and extended relational responding (Barnes-Holmes, Barnes-Holmes, & Cullinan, 2001). According to the REC model, therefore, the data from Experiment 2 indicate that frames of coordination (i.e., the verbal relation of equivalence or similarity) between "Safe" and "White" and between "Dangerous" and "Black" were relatively strong, but frames of distinction (i.e., the verbal relation of difference) between "Dangerous" and "White" and between "Safe" and "Black" were not (the word *strong* is used here simply to denote a high probability in *immediate* relational responding). The REC model assumes that such differences in relational response strengths may be attributed, at least in part, to the verbal and nonverbal contingencies surrounding racial stereotyping. For example, common verbal practices would typically summarize such stereotyping as "white is good" and "black is bad," rather than "white is *not* bad" and "black is *not* good." In other words, two elements of a relational network may well cohere, as in "X is good" and "X is not bad," but the relative strengths or weaknesses of the two elements will be influenced to some degree by other variables, such as differences in frequency of exposure to the two parts of the network.[7] The current results are therefore readily explained by the REC model, although testing the model systematically will have to await further empirical inquiry.

It is worth noting that the specific stereotyping effect observed in Experiment 2 for the Dangerous–Black but not the Safe–Black trial type was also observed at the individual level. Specifically, 8 of the 19 participants produced a negative (pro-black) score on the Safe–Black trial type but a positive (pro-white) score on the Dangerous–Black trial type, with only 1 participant producing the opposite response pattern; of the remaining 10 participants, 7 produced smaller pro-white or larger pro-black effects on the former than on the latter trial type (data available from the first author upon request). Interestingly, this type of effect has been observed previously in a preliminary IRAP study on homonegativity (Cullen & Barnes-Holmes,

7 Although common verbal practices may strengthen frames of coordination over distinction within some relational networks, there are numerous counterexamples. In one recent study (Dawson, Barnes-Holmes, Gresswell, & Giles, 2009), for instance, participants responded far more rapidly to Children–Sexual–False than to Children–Sexual–True. (Interestingly, a group of previously convicted child sex-offenders produced a near zero IRAP effect, responding "True" and "False" with equal speed.) It seems likely, therefore, that cultural taboos may produce relatively strong frames of distinction on the IRAP for nontransgressors because taboos typically involve clear verbal directives on what is *not* permitted.

2008). Such effects could be seen as consistent with recent evidence that indicates the influence of a "negativity bias" in attitude formation (cf. Kunda, 1999). That is, when negatively valenced stimuli are presented with "Black" or "Gay," this serves to activate an implicit anti-black or anti-gay bias, respectively, which is not observed when positively valenced stimuli are presented. On balance, procedural variables specific to the IRAP may be involved here. For example, the stereotyping effect for the Dangerous–Black trial type required responding "True" more quickly than "False," but the opposite was required for the Safe–Black trial type. It is possible, therefore, that a bias toward responding "True" over "False," per se, interacted with the socially loaded stimulus relations presented in the IRAP. If such a response bias does play a role, however, the source of that bias needs to be explained. As suggested previously, the impact of common verbal practices, which tend to confirm negative rather than deny positive stereotypes, is a possibly important variable.

In closing, it is worth noting that the REC model bears some similarity to the associative-propositional evaluation (APE) model, which has been used successfully to explain a wide range of findings in implicit attitudes research (e.g., Gawronski & Bodenhausen, 2007). Unlike the APE model, however, the REC model does not appeal to dual processes (associative and propositional). Rather, the REC model appeals to the single process of arbitrarily applicable relational responding, as defined by RFT.[8] Thus, the divergence between implicit and explicit attitudes is explained not by the complex interplay of factors that affect *separate* associative and propositional processes, but by the extent to which relational responses are elaborated and cohere with each other. Furthermore, the REC model predicts that implicit attitude effects are not restricted to simple associations, but may emerge based on a variety of stimulus relations (see Power et al., 2009, for supporting evidence; see also Deutsch, Kordts-Freudinger, Gawronski, & Strack, 2009). Only further study, however, will determine if the REC model offers clear advantages over the APE model, and our research group is currently engaged in this work.

References

BARNES-HOLMES, D., BARNES-HOLMES, Y., POWER, P., HAYDEN, E., MILNE, R., & STEWART, I. (2006). Do you really know what you believe? Developing the Implicit Relational Assessment Procedure (IRAP) as a direct measure of implicit beliefs. *The Irish Psychologist, 32,* 169–177.

BARNES-HOLMES, D., BARNES-HOLMES, Y., STEWART, I., & BOLES, S. (in press). A sketch of the Implicit Relational Assessment Procedure (IRAP) and the Relational Elaboration and Coherence (REC) Model. *The Psychological Record.*

8 The REC model is *not* a "single process" model because, as a behavior-analytic account, it allows for the involvement of other behavioral processes apart from relational framing, such as respondent conditioning and primary stimulus generalization. Strictly speaking, therefore, the REC model is a *multi*-process model, but one in which the difference between implicit and explicit attitudes is *not* explained by the involvement of distinct behavioral processes; rather, it is the elaboration and coherence involved in the single process of relational framing that provides the core explanation.

BARNES-HOLMES, D., HAYDEN, E., BARNES-HOLMES, Y., & STEWART, I. (2008). The Implicit Relational Assessment Procedure (IRAP) as a response-time and event-related-potentials methodology for testing natural verbal relations: A preliminary study. *The Psychological Record, 58,* 497–516.

BARNES-HOLMES, D., HAYES, S. C., & DYMOND, S. (2001). Self and self-directed rules. In S. C. Hayes, D. Barnes-Holmes, & B. Roche (Eds.), *Relational frame theory: A post-Skinnerian account of human language and cognition* (pp. 119–140). New York: Kluwer/Plenum.

BARNES-HOLMES, D., MURTAGH, L., BARNES-HOLMES, Y., & STEWART, I. (in press). Using the Implicit Association Test and the Implicit Relational Assessment Procedure to measure attitudes toward meat and vegetables in vegetarians and meat eaters. *The Psychological Record.*

BARNES-HOLMES, D., WALDRON, D., BARNES-HOLMES, Y., & STEWART, I. (2009). Testing the validity of the Implicit Relational Assessment Procedure (IRAP) and the Implicit Association Test (IAT): Measuring attitudes towards Dublin and country life in Ireland. *The Psychological Record, 59,* 389–406.

BARNES-HOLMES, Y., BARNES-HOLMES, D., & CULLINAN, V. (2001). Education. In S. C. Hayes, D. Barnes-Holmes, & B. Roche (Eds.), *Relational frame theory: A post-Skinnerian account of human language and cognition* (pp. 181–195). New York: Kluwer/Plenum.

BLANCHARD, F. A., CRANDALL, C. S., BRIGHAM, J. C., & VAUGHN, L. A. (1994). Condemning and condoning racism: A social context approach to interracial settings. *Journal of Applied Psychology, 79*, 993–997.

BOYSEN, G. A., VOGEL, D. L., & MADON, S. (2006). A public versus private administration of the Implicit Association Test. *European Journal of Social Psychology, 36,* 845–856.

BRIEF, A. P., DIETZ, J., COHEN, R. R., PUGH, S. D., & VASLOW, J. B. (2000). Just doing business: Modern racism and obedience to authority as explanations for employment discrimination. *Organizational Behaviour and Human Decision Processes, 81,* 72–97.

CULLEN, C., & BARNES-HOLMES, D. (2008). Implicit pride and prejudice: A heterosexual phenomenon? In T. G. Morrison & M. A. Morrison (Eds.), *The Psychology of Modern Prejudice* (pp. 195–223). New York: Nova Science.

CUNNINGHAM, W. A., PREACHER, K. J., & BANAJI, M. R. (2001). Implicit attitude measures: Consistency, stability, and convergent validity. *Psychological Science, 12,* 163–170.

DASGUPTA, N., MCGHEE, D. E., & GREENWALD, A. G. (2000). Automatic preference for white Americans: Eliminating the familiarity explanation. *Journal of Experimental Social Psychology 36,* 316–328.

DAWSON, D. L., BARNES-HOLMES, D., GRESSWELL, D. M., & GILES, D. (2009). The Implicit Relational Assessment Procedure: Assessing the implicit beliefs of sexual offenders. *Sexual Abuse: A Journal of Research and Treatment, 21*(1), 57-75.

DE HOUWER, J. (2002). The Implicit Association Test as a tool for studying dysfunctional associations in psychopathology: Strengths and limitations. *Journal of Behaviour Therapy and Experimental Psychiatry, 33,* 115–133.

DE HOUWER, J. (2003). The Extrinsic Affective Simon Task. *Experimental Psychology, 50,* 77–85.

DEUTSCH, R., KORDTS-FREUDINGER, R., GAWRONSKI, B., & STRACK, F. (2009). Fast and fragile: A new look at the automaticity of negation processing. *Experimental Psychology, 56(6),* 434–446.

FAZIO, R. H., SANBONMATSU, D. M., POWELL, M. C., & KARDES, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology, 50,* 229–238.

GAWRONSKI, B., & BODENHAUSEN, G. V. (2007). Unraveling the processes underlying evaluation: Attitudes from the perspective of the APE model. *Social Cognition, 25*, 687–717.

GAWRONSKI, B., LEBEL, E. P., & PETERS, K. R. (2007). What do implicit measures tell us? Scrutinizing the validity of three common assumptions. *Perspectives on Psychological Science, 2,* 181–193.

GREENWALD, A. G., BANAJI, M. R., RUDMAN, L. A., FARNHAM, S. D., NOSEK, B. A., & MELLOTT, D. S. (2002). A unified theory of implicit attitudes, stereotypes, self-esteem, and self-concept. *Psychological Review, 109,* 3–25.

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., OAKES, M. A., & HOFFMAN, H. G. (2003). Targets of discrimination: Effects of race on responses to weapons holders. *Journal of Experimental Social Psychology, 39,* 399–405.

HAYES, S. C., BARNES-HOLMES, D., & ROCHE, B. (2001). *Relational frame theory: A post-Skinnerian account of human language and cognition*. New York: Kluwer/Plenum.

KUNDA, Z. (1999). *Social cognition: Making sense of people*. Cambridge, MA: MIT Press.

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.

LOWERY, B. S., HARDIN, C. D., & SINCLAIR, S. (2001). Social influence effects on automatic racial prejudice. *Journal of Personality and Social Psychology, 81*, 842–855.

MAASS, A., CASTELLI, L., & ARCURI, L. (2000). (2000). Measuring prejudice: Implicit versus explicit techniques. In D. Capozza & R. Brown (Eds.), *Social identity processes: Trends in theory and research* (pp. 96–116). London: Sage.

MCCONAHAY, J. B. (1986). Modern racism, ambivalence, and the Modern Racism Scale. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 91–125). Orlando, FL: Academic Press.

MCKENNA, I. M., BARNES-HOLMES, D., BARNES-HOLMES, Y., & STEWART, I. (2007). Testing the fake-ability of the Implicit Relational Assessment Procedure (IRAP): The first study. *International Journal of Psychology and Psychological Therapy, 7*, 253–268.

MONTEITH, M. J., VOILS, C. I., & ASHBURN-NARDO, L. (2001). Taking a look underground: Detecting, interpreting, and reacting to implicit racial biases. *Social Cognition, 19,* 395–417.

MOORS, A., & DE HOUWER, J. (2006). Automaticity: A theoretical and conceptual analysis. *Psychological Bulletin, 132,* 297–326.

NOSEK, B., & BANAJI, M. R. (2001). The go/no-go association task. *Social Cognition, 19,* 625–666.

NOSEK, B., GREENWALD, A. G., & BANAJI, M. R. (2004). Understanding and using the Implicit Association Test: II. Method variables and construct validity. *Personality and Social Psychology Bulletin, 31,* 166–180.

OTTAWAY, S. A., HAYDEN, D. C., & OAKES, M. A. (2001). Implicit attitudes and racism: The effect of word familiarity and frequency in the Implicit Association Test. *Social Cognition, 19,* 97–144.

PLANT, E. A., & DEVINE, P. G. (1998). Internal and external motivation to respond without prejudice. *Journal of Personality and Social Psychology, 75,* 811–832.

POWER, P. M., BARNES-HOLMES, D., BARNES-HOLMES, Y., & STEWART, I. (2009). The Implicit Relational Assessment Procedure (IRAP) as a measure of implicit relative preferences: A first study. *The Psychological Record, 59*, 621–640.

PRATTO, F., & JOHN, O. P. (1991). Automatic vigilance: The attention grabbing power of negative social information. *Journal of Personality and Social Psychology, 61,* 380–391.

RICHESON, J. A., & AMBADY, N. (2003). Effects of situational power on automatic racial prejudice. *Journal of Experimental Social Psychology, 39,* 177–183.

VAHEY, N. A., BARNES-HOLMES, D., BARNES-HOLMES, Y., & STEWART, I. (2009). A first test of the Implicit Relational Assessment Procedure as a measure of self-esteem: Irish prisoner groups and university students. *The Psychological Record, 59*, 371–388.

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.

The Psychological Record

The Psychological Record (ISSN 0033-2933) is a nonprofit journal of psychology appearing quarterly. Subscriptions are available on a calendar year basis (January through December) and it is published by Southern Illinois University Carbondale. *The Psychological Record* is primarily for the publication of experimental or theoretical papers in many areas of psychology. The journal was founded by renowned interbehaviorist J.R. Kantor, with one of the most influential American psychologists, B.F. Skinner, serving as the experimental department editor at the journal's inception.

DOES YOUR LIBRARY SUBSCRIBE TO THE PSYCHOLOGICAL RECORD?

IF NOT, THEN COPY THIS FORM FOR YOUR REFERENCE LIBRARIAN.

To be filled out by employee or faculty member
I would like to request that our library purchase a 2010 subscription to
The Psychological Record (ISSN 0033-2933)

Name _____ Department _____

Signature _____

To be filled out by the reference librariran
The Psychological Record (ISSN 0033-2933) Volume 60 (2010):
$160 (USA) $170 (International)

Institution Name _____

Address _____

City _____ State _____ Zip _____ Country _____

Payment: Check (in US currency) made payable to: The Psychological Record

VISA MasterCard Credit Card #: _____

Expires: _____

Security Code _____

Signature: _____

Mail to: The Psychological Record
Rehabilitation Institute
Rehn Hall - Mail Code 4609
Southern Illinois University Carbondale
1025 Lincoln Drive
Carbondale IL 62901