Anti-fat, pro-slim, or both?

Using two reaction-time based measures to assess implicit attitudes to the slim and overweight

SARAH RODDY & IAN STEWART
National University of Ireland, Galway
DERMOT BARNES-HOLMES
National University of Ireland, Maynooth

Abstract

Two measures of implicit attitudes, the Implicit Relational Assessment Procedure (IRAP) and the Implicit Association Test (IAT), were compared with each other and with a measure of explicit attitudes in the assessment of implicit pro-slim/anti-fat bias. Results from both implicit tests indicated higher levels of bias than revealed by the explicit measure. The IRAP data suggested that it was participants' pro-slim rather than anti-fat bias, which was driving this effect. Explicit attitudes and feelings towards the overweight were significant predictors of behavioural intentions towards the overweight with the IRAP offering a greater contribution to predictive validity than the IAT.

Recently, implicit measures have been employed in research examining anti-fat bias (e.g. Bessenoff & Sherman, 2000; Teachman, Gapinski, Brownell, Rawlins, & Yeyaram, 2003). The most prominent such measure is the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998). Research into anti-fat bias using the IAT typically finds that response latencies are faster when 'slim' stimuli and positive words (e.g. 'desirable') and 'fat' stimuli and negative words (e.g. 'undesirable') must be categorised together than when 'slim' stimuli and negative words and 'fat' stimuli and positive words must be categorised together. Furthermore, this anti-fat IAT effect has been replicated in numerous studies using both textual and pictorial stimuli (see Morrison, Roddy, & Ryan, 2009).

A consistent trend has emerged such that participants produce significantly higher levels of bias on the IAT than on explicit measures (e.g., Schwartz, Vartanian, Nosek, & Brownell, 2006; Teachman et al., 2003). Thus, the IAT appears to reveal levels of anti-fat prejudice that might not otherwise come to light. Despite its popularity, however, the IAT possesses several limitations. One consistent criticism is that the associations found for any particular concept are always relative (De Houwer, 2002). For example, an IAT effect for 'slim' over 'fat' could indicate a neutral attitude towards slim people and a negative attitude towards fat people, or a positive attitude towards slim people and a neutral attitude towards fat people. Hence, the standard IAT does not provide a measure of the separate attitudes to 'fat' and 'slim'.

To date, one study employed a priming procedure to assess separate implicit attitudes in the context of body weight (Bessenoff & Sherman, 2000). Participants were presented with photographs of overweight and thin women, followed by a lexical decision task (LDT) in which they had to make judgments about words that were negatively or positively valenced. The dependent measure was response latency in making word/non-word judgments. Additionally, these researchers assessed behavioural intention towards an overweight woman by examining preferred seating distance. Physical distance between a chair that participants were led to believe an overweight woman was to sit on, and a chair that participants were instructed to place in the room, was measured, with greater seating distance indicating more negative behavioural intent.

Overall, there was greater activation for negative trait words when preceded by the photo of the overweight than the thin models. Activation of the positive trait words was not greater when preceded by either type of photos. These data indicated that the overweight models were perceived negatively, but the slim models were perceived in a relatively neutral manner. Automatic evaluations were found to correlate with behavioural intention, such that participants evidencing greater implicit negativity chose to sit further away from the overweight woman. Explicit attitudes did not correlate with behavioural intention suggesting that the LDT had greater predictive validity.

The current study sought to examine implicit attitudes to the overweight and slim using both the IAT and a new *non*-relative measure, the Implicit Relational Assessment Procedure (IRAP; Barnes-Holmes et al., 2006). Similar to the IAT, the IRAP is a computer-based task in which participants must respond either consistently or inconsistently with hypothesised implicit attitudes, and the difference in response latency between consistent and inconsistent trials provides the critical metric. Unlike the IAT, however, the IRAP requires that participants respond to four different stimulus relations rather than simply categorizing stimuli into two groups.

To assess attitudes to body-weight using an IRAP, participants might be presented with trials in which a positive word (e.g. 'Good') or a negative word (e.g. 'Bad') is presented at the top of a computer screen with a picture of an overweight or slim individual presented in the middle of the screen; on each trial participants may be required to choose between the response options, 'Same' and 'Opposite'. Four different trial-types would thus be created by presenting each of the two words with each of the two types of pictures; *Good-Slim*, *Bad-Slim*, *Good-Overweight*, *Bad-Overweight.* During such an IRAP, two types of blocks of trials would be presented with one type requiring responses that are deemed pro-slim and anti-fat (i.e. *Good-Slim-Same*; *Bad-Slim-Opposite*; *Good-Overweight-Opposite*; *Bad-Overweight-Same*), and another requiring a pro-fat/anti-slim response pattern (e.g. *Good-Slim-Opposite*, *Good-Overweight-Same*, etc.). The IRAP requires that participants alternate between these two types of blocks while responding as accurately and as quickly as possible. Critically, four separate latency difference scores, one for each trial-type, are yielded by the IRAP. For example, if participants respond more quickly to *Good-Slim-Same* than to *Good-Slim-Opposite* this would indicate a pro-slim bias; if participants respond more quickly to

Bad-Overweight-Same than to *Bad-Overweight-Opposite* this would indicate an anti-fat bias.

Although the IRAP is relatively new, a few studies have shown that it: (1) 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); (2) is not easily faked (McKenna, Barnes-Holmes, Barnes-Holmes, & Stewart, 2007); (3) may be used as a measure of implicit self-esteem (Vahey, Barnes-Holmes, Barnes-Holmes, & Stewart, 2009); and (4) produces effects that clearly diverge from those obtained from explicit measures when targeting socially sensitive attitudes (Barnes-Holmes, Murphy, Barnes-Holmes, Stewart, in press; Power, Barnes-Holmes, Barnes-Holmes, & Stewart, 2009). At the current time, however, no published study has attempted to use the IRAP to assess implicit attitudes to the slim and overweight.

As noted previously, only one article has reported the non-relative assessment of implicit attitudes to body-weight (Bessenoff & Sherman, 2000), and the data indicated a negative-overweight bias but a neutral effect for thin. Based on this finding one might predict a similar outcome for the IRAP. In the current study, however, pictures of average-weight men and women were presented as 'thin' stimuli, rather than the computer-adjusted thin female models used by Bessenoff and Sherman. Consequently, we refrained from making specific predictions concerning individual trial-type effects for the IRAP. However, consistent with previous research we did predict an IAT effect, and an overall IRAP effect (i.e. averaged across the four trial-types), indicative of a relative implicit bias that favored the slim/average-weight photographs.

The current study also sought to explore the possible relation between implicit attitudes and behavioural variables. To date, there have been three IAT investigations of the relationship between implicit anti-fat bias and behavioural intentions towards the overweight (e.g. Brochu & Morrison, 2007; Gapinksi, Schwartz, & Brownell, 2006; O'Brien et al., 2008). However, inconsistent results have emerged in these studies. Specifically, Gapinksi et al. (2006) reported that IAT and explicit attitudinal measures were related to willingness to have an overweight target as a roommate. However, Brochu and Morrison found that only the explicit measure was related to behavioural intentions towards an overweight target. Furthermore, a recent study found

that neither the IAT nor the explicit measure of anti-fat attitudes predicted job-related discrimination measures (O'Brien et al., 2008). The issue is further complicated by the fact that the priming study by Bessenoff and Sherman (2000) indicated that implicit attitudes were predictive of behavioural intentions. Given the limited and contradictory evidence in this area, making a specific prediction was difficult and thus we sought simply to determine if the IAT and/or the IRAP predicted behavioural intentions towards pictures of overweight versus average-sized females.

Method

Participants

Eighty participants, 58 females and 22 males, aged 18 to 40 years ($M = 21$ years and 1 month, $SD = 3$ years and 4 months) were randomly assigned to four conditions (see later). Participants were undergraduate and postgraduate psychology students at a university within the Republic of Ireland, none of whom had previous exposure to implicit measures.

Materials and stimuli

Instructional and stimulus presentation and the recording of participants' responses for the implicit measures were conducted using a standard laptop computer.

IAT

The stimuli employed in the IAT task included 12 images (6 photos of males and 6 of females, with half of each gender overweight and half average weight) and 6 positive (desirable, active, disciplined, attractive, healthy, good) and 6 negative (undesirable, lazy, sloppy, ugly, ill, bad) adjectives. Images were selected from a series of 24 photographic stimuli employed by Brochu and Morrison (2007), which had been pilot tested for target gender and weight status representativeness and for which slim and fat photographs were matched for levels of attractiveness.

IRAP

Participants were presented with the adjectives 'Good' and 'Bad' as label stimuli. The set of 12 images employed in the IAT were also used as target stimuli for the IRAP. Additionally, participants were presented with the two relational response options, 'Same' and 'Opposite'.

General Information Questionnaire (GIQ)

This was a 26-item questionnaire designed specifically for the purposes of this research and employed to assess relevant participant information (e.g. weight, height, amount of contact with overweight individuals on an everyday basis, previous exposure to the implicit measures). Two items asked participants to rate how warm or cold they felt towards overweight people and slim people; scores ranged from 1 to 5 with low scores indicating warmer feelings (1 = *very warm*) and high scores indicating colder feelings (5 = *very cold*).

Explicit Anti-Fat Attitudes (AFA)

The AFA (Crandall, 1994) was used to assess participants' explicit attitudes towards overweight individuals. This scale contains 13-items representing three subscales: (1) dislike of fat people (7 items); (2) personal fear of fat (3 items); and (3) belief in the willpower of overweight individuals (3 items). Consistent with previous research conducted by O'Brien, Hunter and Banks (2006) the word 'weight' was removed from one of the items in the fear of fat subscale and replaced with the word 'fat'. Participants rated their agreement with the statements on a 10-point rating scale (0 = True; 9 = False); therefore scores ranged from 0 to 9 and were scored so that higher scores indicated greater endorsement of the item (i.e. greater anti-fat bias). The AFA showed satisfactory scale score reliability. Cronbach's alpha for each of the subscales was very good; .80 (fear of fat; 95% CI = .70–.88),.74 (dislike; 95% CI = .62–.83) and .74 (willpower; 95% CI =.60–.84). Cronbach's Alpha for the overall scale was .67 (95% CI = .53–.78).

Behavioural Intention Questionnaire (BIQ)

Based on work by Brochu and Morrison (2007) participants were presented with two photographs; one depicting an overweight female, the other an average weight female. For each of the two photos, the participant had to answer five questions assessing the extent to which they would interact with the pictured person. Each question involved a 7-point rating scale (1 = *very unlikely* to 7 = *very likely*). Therefore, scores could range from 5 to 35, with higher scores indicating greater likelihood of interaction with the target. Cronbach's alpha for the scales were .87 (average weight; 95% CI = .82 –.92) and .92 (overweight; 95% CI = .88 –.95), respectively.

Procedure

Ethical approval for the study was granted by the Research Ethics Committee at the host university. The experiment consisted of two phases. Phase 1 involved exposure to both implicit tests while Phase 2 involved completion of the explicit measures. The order in which the two implicit measures were presented was counterbalanced across participants. During Phase 1, participants were provided with a 4–5 minute break after completing the first implicit test and before exposure to the second. Participants were presented with implicit tests that commenced with blocks that required a pro-slim/anti-fat response pattern or the opposite (anti-slim/pro-fat), and this variable was counterbalanced across participants. Thus, four conditions were created: (1) IAT->IRAP/pro-slim/anti-fat; (2) IAT->IRAP/anti-slim/pro-fat; (3) IRAP->IAT/pro-slim/anti-fat; and, (4) IRAP->IAT/anti-slim/pro-fat.

Phase 1: The implicit protocols

The IAT Participants were seated in front of the computer, which provided a detailed set of task instructions. They were asked to read the instructions and to consult the experimenter if they were uncertain of how to respond.

The IAT that commenced with a pro-slim/anti-fat response sequence was as follows. There were seven blocks, starting with a practice block requiring the categorization of stimuli related to the categories 'Slim Person' and 'Fat Person'. On each trial a photographic stimulus was presented center screen and the participant was required to categorize it by pressing either the 'd' or 'k' key. If categorized correctly, the stimulus was removed and the next stimulus was presented 400ms later. If categorized incorrectly, a red 'X' appeared below the stimulus, and the participant had to select the appropriate response key. This block consisted of 24 trials, with the 12 target stimuli presented in a quasi-random order and each appearing twice. After the 24th trial, participants were presented with feedback concerning their percentage of correct responses and the median response latency for that block.

The procedure for Block 2 was similar to Block 1, except that participants were required to categorize the positive and negative adjectives into the categories 'Good' and 'Bad'. In Block 3, participants had to categorize the stimuli as 'Good or Slim Person' or 'Bad or Fat Person'. The 12 photographs and 12 adjectives were presented quasi-randomly with the constraint that each of the 24 stimuli appeared once across the 24 trials. Block 4 was similar to Block 3 except that it was

a test block, and it consisted of 48 trials, with each target presented twice in a random order. Block 5 was similar to Block 1, except that the category labels switched locations from one side of the screen to the other. Block 6 was a practice block similar to Block 3, except that stimuli had to be categorized into 'Good or Fat Person' and 'Bad or Slim Person' categories. Block 7 was similar to Block 6, except that it was a test block and consisted of 48 trials, with each target stimulus presented twice in a quasi-random order.

The procedure for the IAT that commenced with a pro-fat/anti-slim response sequence was similar to that described above, except that participants were exposed to the blocks in the following order: Block 5, Block 2, Block 6, Block 7, Block 1, Block 3, and Block 4.

The IRAP The IRAP consisted of a minimum of two practice blocks and six test blocks. Two IRAP sequences were employed – pro-slim/anti-fat-first and pro-fat/anti-slim-first. On each trial of the first practice block of the pro-slim/anti-fat-first sequence the label stimulus (either 'Good' or 'Bad') was presented at the top centre of the screen. A target stimulus (one of the 12 photographic stimuli) was presented directly below the label, in the middle of the screen. Finally, the two relational terms, 'Same' and 'Opposite' appeared at the bottom of the screen, one on the left and one on the right. Participants were required to 'indicate' the relation between the label and target stimuli by choosing either 'Same' or 'Opposite' by pressing either the 'd' (left) or 'k' (right) key. A correct response initiated a 400 ms inter-trial interval followed by the next trial, whereas an incorrect response produced a red 'X' mid-screen, after which the participant had to make the correct response before continuing to the inter-trial interval.

Each block consisted of 24 trials with each of the 12 photographic stimuli presented twice in a quasi-random order, once in the presence of each of the label stimuli ('Good' and 'Bad'). In the first practise block, participants were required to respond in accordance with a pro-slim/anti-fat bias; for example, given 'Good' and an 'average weight' photo the relational term 'Same' was correct, whereas given 'Bad' and the 'average weight' photo the relational term 'Opposite' was correct. On completion of the block, participants were provided with feedback detailing percentage of correct responses and median response latency. They then received instructions informing them that the previously correct and incorrect answers would be reversed in the next block.

Practise Block 2 was similar to Block 1, except that the participants were required to respond in accordance with a pro-fat/anti-slim bias. Thus, for example, given 'Good' and the 'average weight' photo the relational term 'Opposite' was correct, whereas given 'Bad' and the 'average weight' photo the relational term 'Same' was correct.

If participants achieved 80 per cent correct and a median response latency of 3000 milliseconds or less on each of the practice blocks, then they advanced to the 6 test blocks. Any participant who failed to reach the practise criteria was re-exposed to practise Blocks 1 and 2 up to a maximum of four cycles (i.e. a total of eight blocks). Participants who failed to reach criterion within four exposures were thanked, debriefed and their data were discarded. A total of 64 out of 80 participants achieved criterion (and, in addition, maintained 80% accuracy across the six IRAP test blocks) and their data was retained for analysis.

The procedure for the test blocks was similar to that for the practise blocks, except that an instruction was presented on-screen before each block stating that 'this is a test' and to 'go fast' while also trying to make few errors. No performance criteria were required to continue from one test block to the next.

The pro-fat/anti-slim-first sequence was similar to that described earlier except that the IRAP commenced with a practise block that required participants to respond in accordance with a pro-fat/anti-slim bias; all subsequent blocks, practise and test, then alternated between pro-slim/anti-fat and pro-fat/anti-slim.

Phase 2: Explicit attitudinal measures

On completion of the second implicit protocol, participants completed the GIQ, AFA, and BIQ attitudinal measures and were then thanked and debriefed.

Results

A series of preliminary analyses showed that neither order of administration of the implicit tests (IAT > IRAP or IRAP > IAT) nor block sequence (pro-slim/anti-fat-first or pro-fat/anti-slim-first) had a significant effect on performance (all $ps > .51$); hence these variables were removed from subsequent analyses.

IAT

The primary datum in the case of the IAT was response latency defined as time in milliseconds from the onset of a trial until the emission of a correct response.

The response latency data for each participant were transformed into *D*-scores using the *D*-algorithm, developed by Greenwald, Nosek, and Banaji (2003), which controls for individual differences in response speed. For the current study, the latencies from pro-slim/anti-fat blocks were subtracted from pro-fat/anti-slim blocks, and thus a positive *D*-score indicates that a participant responded more quickly when categorizing positive adjectives with slim and negative with fat, than when categorizing in the opposite manner (slim with negative and fat with positive). The average *D*-IAT score across participants was .68 (*SD* = .42) and significantly different from zero (*t* [63] = 13.1, *p* < .001, *r* = .86), thus yielding the predicted pro-slim/anti-fat bias.

IRAP

Similar to the IAT, the primary datum for the IRAP was response latency defined as time in milliseconds from the onset of a trial until the emission of a correct response. Individual response latency data were transformed into *D*-IRAP scores (e.g. Barnes-Holmes, Murtagh, et al., in press) using an adaptation of the Greenwald et al. (2003) *D*-algorithm. As for the IAT, the latencies from pro-slim/anti-fat blocks were subtracted from pro-fat/anti-slim blocks, and thus a positive *D*-IRAP score indicates a pro-thin or anti-fat bias, whereas a negative score indicates the opposite. The overall *D*-IRAP scores for each of the four trial types are presented in Fig. 1 (the overall mean *D*-IRAP score was .14 [*SD* = .29], and significantly different from zero, *t* [63] = 3.81, *p* < .001, *r* = .43). A relatively strong pro-slim IRAP effect was observed for the *Good-Slim* trial-type (responding 'Same' faster than 'Opposite') and for the *Bad-Slim* trial-type (responding 'Opposite' faster than 'Same'). In contrast, relatively weak pro-fat IRAP effects were observed for *Good-Fat* (responding 'Same' faster than 'Opposite') and *Bad-Fat* (responding 'Opposite' faster than 'Same') trial-types. One sample *t*-tests indicated that the *D*-IRAP scores for the two *Slim* trial types were statistically significant from zero (*Good-Slim*, *t* [63] = 5.47, *p* < .001, *r* = .57; *Bad-Slim*, *t* [63] = 5.31, *p* < .001, *r* = .56), whereas the scores for the two *Fat* trial-types were not statistically significant. The *Bad-Fat* trial-type was significantly correlated with both the *Bad-Slim*, *r* (62) = .30, *p* = .017, and the *Good-Fat*, *r* (62) = .25, *p* = .045, trial-types. No other statistically significant relationships emerged between the trial-types (*rs* < .24, *ps* > .061). Overall, therefore, the IRAP indicated a positive implicit bias towards the average-weight photographs, with the absence of any
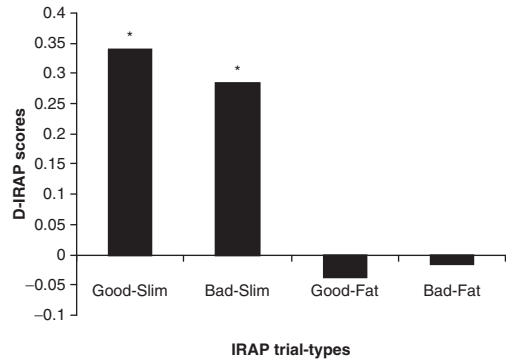

Figure 1. D-IRAP scores for the four trial-types
**p*<.05*

bias towards the photographs of overweight individuals. In short, the participants were pro-slim rather than anti-fat.

Comparison of the IRAP and IAT

The *D*-scores for the IAT and the overall *D*-IRAP scores were not statistically correlated, *r* (62) = .18, *p* = .15, but there was a statistically significant relationship between the *D*-IAT scores and the *Bad-Slim* trial-type, *r* (62) = .27, *p* = .031.

Explicit anti-fat attitudes: The AFA scale

Average scores on the AFA subscales were 1.82 (*SD* = 1.54) for dislike, 5.16 *(SD* = 1.95) for willpower, and 4.35 (*SD* = 2.53) for fear of fat, in all cases scores ranged from 0 to 9. Therefore, participants reported relatively low levels of explicit anti-fat attitudes.

GIQ

Participants expressed warm attitudes towards both overweight people (*M* = 2.77, *SD* = .66) and slim people (*M* = 3.39, *SD* = .542), with a statistically significant difference in ratings *t* (61) = -4.76, *p* < .001, *r* = -.46; feelings towards overweight people were more positive than feelings towards normal weight individuals.

GIQ – AFA correlations

A statistically significant relationship emerged between feelings towards the overweight and dislike on the AFA, *r* (56) =.33, *p* = .010.

BIQ

The five items of the BIQ were summated creating a total scale score. More positive behavioural intentions

were recorded towards the average weight (M = 23.06; SD = 4.89) than the overweight target (M = 20.94; SD = 5.44), and this difference proved to be statistically significant, t (56) = 2.40, p = .020, r = .50.

Implicit–explicit correlations

No statistically significant relationships emerged between the three subscales of the AFA, or the feelings towards the overweight scale and either of the implicit tests (rs < .25, ps > .058). Three correlations yielded ps < .1; overall D-IRAP score and dislike scores, r (62) = .25, p = .058; dislike scores and the *Bad-Fat* and *Good-Fat* trial types, r (62) = .245, p = .063 and r (62) = .239, p = .071, respectively.

Predictive validity

Regression analyses were conducted to determine if the IRAP and the IAT would increase the prediction of behavioural intentions towards an overweight target offered by explicit dislike and explicit feelings towards the overweight. For all variables, data met the principal statistical assumptions required for multiple regression analyses.

The AFA

For all models the dislike subscale of the AFA was entered in the first step and proved to be a statistically significant predictor of behavioural intentions, F (1,52) = 9.14, β = −.39, p = .004, adjusted R^2 = .13. In the first model, the D-IAT score was entered as step 2, and this produced virtually no increment in predictive validity, β = −.04, p = .74, R^2 change = .00. In the second model, the overall D-IRAP score was entered into step 2 and although not statistically significant, β = −.15, p = .25, R^2 change = .02, it produced a larger increment in the prediction of behavioural intention than offered by the IAT.

Feelings towards the overweight

For all models, feelings towards the overweight was entered as step 1 and proved to be a statistically significant predictor of behavioural intentions, F(1, 55) = 8.92, β = −.368, p = .004, adjusted R^2 = .12. In the first model, the D-IAT score was added in step 2 and this produced virtually no increment in the predictive validity of the model, β = −.03, p = .79, R^2 change = .00. In the second model, the overall D-IRAP score was entered into step 2, β = −.24, p = .055, R^2 change = .057, adding 5.7 per cent to the prediction of behavioural intentions towards the overweight target.

Discussion

Consistent with previous IAT-based research, Irish participants demonstrated strong levels of implicit pro-slim/anti-fat bias, which was not observed with the AFA (e.g. Schwartz et al., 2006; Teachman & Brownell, 2001). Similar to this, the IRAP also showed an overall effect indicative of anti-fat/pro-slim bias. However, the non-relative nature of the IRAP revealed participants demonstrated pro-slim attitudes towards the average-weight stimuli rather than anti-fat attitudes towards the over-weight photographs. In effect, the IRAP showed that participants were pro-slim rather than anti-fat. The implicit measures did not appear to be strongly related, in that only the *Bad-Slim* trial-type correlated with the IAT. A non-significant trend suggested that the overall D-IRAP and the *Good-Fat* and *Bad-Fat* trial-types related to explicit dislike towards the overweight. Consistent with previous IAT research, this measure failed to relate to explicit anti-fat attitudes (e.g. Ahern & Hetherington, 2006). Only the IRAP increased predictive validity with respect to behavioural intentions. It is important to note, however, that restriction in range may be problematic, possibly resulting in the low correlations evidenced across the implicit and explicit measures (Zimmerman & Williams, 2000).

The current data provide additional support for the utility of explicit measures in predicting behavioural intentions towards overweight targets. Both Brochu and Morrison (2007) and Gapinski et al. (2006) reported similar results in their investigations. In addition, the current study provides initial evidence for the role of explicit feelings towards the overweight in predicting behavioural intentions with both feelings of warmth and low levels of dislike significantly predicting less negative behavioural intentions towards the overweight female target. The fact that participants reported significantly greater warmth towards overweight than average weight people and relatively low levels of dislike could, of course, reflect socially desirable responding (Brochu & Morrison, 2007).

As discussed earlier, one published study employed a non-relative measure of implicit attitudes to body-weight, and the findings showed that overweight models were perceived negatively, whereas the slim models produced a neutral effect (Bessenoff & Sherman, 2000). The IRAP produced the opposite pattern (pro-slim and neutral-fat) and thus future research will need to determine if the divergent

outcomes resulted from the different types of stimuli that were employed across the two studies (i.e. exclusively female versus male and female; and computer-adjusted thin models versus average-weight individuals). In any case, the current findings serve to highlight that the anti-fat bias observed in previous IAT studies may be driven, at least in part, by a pro-slim effect. Indeed, this outcome may not be surprising in that the thin ideal is ubiquitous in the media and those who obtain and maintain it are placed on a social pedestal (Bordo, 1993).

Both explicit dislike and feelings towards the overweight emerged as significant predictors of behavioural intentions towards the overweight target. The finding that the IAT did not contribute to the prediction of behavioural intentions is consistent with results of other recent studies (e.g. Brochu & Morrison, 2007; O'Brien et al., 2008). Interestingly, in comparison with the IAT the IRAP did produce increases in the predictive validity of the explicit measures, with one increase reaching marginal significance. The current data thus suggest a trend that favours the IRAP over the IAT in predicting participants' behavioural intentions towards an overweight target. Although the reason behind this outcome remains unclear at the present time, it is worth noting that the only other study that used a non-relative measure in this area also reported that it predicted behavioural intentions (Bessenoff & Sherman, 2000). Clearly further work is needed, to understand how different implicit measures either predict or fail to predict behaviour.

To the authors' knowledge, the current study represents the first empirical investigation into the anti-fat attitudes of Irish university students. The IAT and explicit data, in particular, reveal patterns of bias that are similar to those reported with samples of participants from America (e.g. Gapinski et al., 2006; Teachman et al., 2003), Canada (e.g. Brochu & Morrison, 2007), Britain (e.g. Ahern & Hetherington, 2006), and New Zealand (e.g. O'Brien et al., 2006, 2008); suggesting the nature of anti-fat attitudes and, particularly, the inappropriateness of their overt expression is also similar in an Irish context.

Specifically with regard to IRAP performance, it should be noted that a substantial number of participants did not reach criterion performance on the practise blocks and were consequently removed from analyses, whereas all participants completed the IAT. On balance, a performance criterion is not typically imposed on the practise blocks of the IAT (in order to progress to the test blocks) and thus it is difficult to compare the two measures directly on this basis. Nevertheless, future research that aims to compare the IRAP with the IAT should remain alert to this perhaps important procedural difference. For example, future studies might impose a performance criterion on the IAT or not employ one with the IRAP. The latter option might prove unwise, however, because anecdotal evidence suggests that some participants find the IRAP a relatively difficult task, and thus they require considerable practise before they are fully able to follow the instructions. It is worth noting that recent evidence indicates that the automaticity or 'implicitness' of the IRAP may be increased by reducing the latency practice criterion from 3000 to 2000 milliseconds (Barnes-Holmes, Murphy, et al., in press). This reduction almost certainly increases the difficulty of the task, thus further bolstering the need for a performance criterion.

One criticism of the current study might be that participants were presented with images of average weight and overweight male and female targets within the implicit tasks, but behavioural intentions were assessed with respect to an overweight female target. Although Brochu and Morrison (2007) adopted a similar approach, participants in their study were presented with *either* male or female targets in the implicit and explicit measures and behavioural intentions were assessed for the corresponding target only. Critically, Brochu and Morrison also found that their IAT did not predict behavioural intentions, and thus it seems unlikely that matching the gender of the stimuli across the implicit and explicit measures in the current study would have yielded different results for the IAT. On balance, it remains unclear what impact greater consistency in this regard would have had on the IRAP measure, and thus future research might explore this issue.

A related concern is that across the implicit and the explicit measures, attitudes were assessed towards 'overweight' people, in general, and the reliance on terms that were not gender-specific may have reduced the ability of these measures to accurately represent bias towards the overweight. Undoubtedly, it is possible that evaluations of 'overweight men' and 'overweight women' may yield disparate results. Given the effectiveness of the IRAP in this domain, an interesting avenue for further research involves independently examining the nature of implicit anti-fat bias towards males and females.

Overall, the current study provides insight into the nature of implicit anti-fat/pro-slim bias. Indeed,

the data suggest several possible advantages of the IRAP over the IAT in this context. In particular, the trial-type data of the IRAP indicate that the 'anti-fat' bias observed in previous IAT studies may be driven, at least in part, by a pro-slim bias. In addition, the emergent trend from the regression analyses suggest, albeit tentatively, that the IRAP may offer a greater contribution than the IAT in increasing the predictive validity of some explicit measures. The findings thus highlight the potential for the IRAP in future investigations of implicit anti-fat/pro-slim bias and in the investigation of the relationship between this bias and behavioural variables.

References

Ahern, A. L., & Hetherington, M. M. (2006). The thin ideal and body image: An experimental study of implicit attitudes. *Psychology of Addictive Behaviour*, *20*, 338–342.

Barnes-Holmes, D., Barnes-Holmes, Y., Hayden, E., Milne, R., Power, P., & Stewart, I. (2006). Do you really know what you believe?: Developing the Implicit Relational Assessment Procedure as a direct measure of implicit beliefs. *The Irish Psychologist*, *32*, 169–177.

Barnes-Homes, D., Murphy, A., Barnes-Holmes, Y., & Stewart, I. (in press). The Implicit Relational Assessment Procedure (IRAP): Exploring the impact of private versus public contexts and the response latency criterion on pro-white and anti-black stereotyping among white Irish individuals. *The Psychological Record*.

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 towards 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: Measuring attitudes towards Dublin and country-life in Ireland. *The Psychological Record*, *59*, 389–406.

Bessenoff, G. R., & Sherman, J. W. (2000). Automatic and controlled components of prejudice towards fat people: Evaluation versus stereotype activation. *Social Cognition*, *18*, 329–353.

Bordo, S. B. (1993). *Unbearable weight: Feminism, western culture, and the body.* Berkley, CA: California University Press.

Brochu, P. M., & Morrison, M. A. (2007). Implicit and explicit prejudice toward overweight and average weight men and women: Testing their correspondence and relation to behavioral intentions. *Journal of Social Psychology*, *147*, 681–706.

Crandall, C. S. (1994). Prejudice against fat people: Ideology and self-interest. *Journal of Personality & Social Psychology*, *66*, 882–894.

De Houwer, J. (2002). The implicit association test as a tool for studying dysfunctional associations in psychopathology: Strengths and limitations. *Journal of Behavior Therapy & Experimental Psychiatry*, *33*, 115–133.

Gapinski, K. D., Schwartz, M. B., & Brownell, K. D. (2006). Can television change anti-fat attitudes and behaviour? *Journal of Applied BioBehavioural Research*, *11*, 1–28.

Greenwald, A. G., Mc Ghee, A. D., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality & Social Psychology*, *74*, 1464–1480.

Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I. An improved scoring algorithm. *Journal of Personality & Social Psychology*, *85*, 198–216.

McKenna, I., 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 & Psychological Therapy, 7,* 253–268.

Morrison, T. G., Roddy, S., & Ryan, T.A. (2009). Methods for measuring attitudes about obese people. In D. B. Allison & M. L. Baskin (Eds.), *Handbook of assessment methods for obesity and eating behaviors,* pp. 79–113. New York: SAGE.

O'Brien, K. S., Hunter, J. A., & Banks, M. (2006). Implicit anti-fat bias in physical educators: Physical attributes, ideology, and socialization. *International Journal of Obesity*, *1*, 308–314.

O'Brien, K. S., Latner, J. D., Halberstadt, J., Hunter, J. A., Anderson, J., & Caputi, P. (2008). Do anti-fat attitudes pedict anti-fat behaviours? *Obesity*, *16*, 87–92.

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.

Schwartz, M. B., Vartanian, L. R., Nosek, B. A., & Brownell, K. D. (2006). The influence of one's own body weight on implicit and explicit anti-fat bias. *Obesity*, *14*, 4404–4447.

Teachman, B. A., & Brownell, K. D. (2001). Implicit anti-fat bias among health professionals: Is anyone immune? *International Journal of Obesity*, *25*, 1525–1531.

Teachman, B. A, Gapinski, K. D., Brownell, K.D., Rawlins, M., & Yeyaram, S. (2003). Demonstrations of implicit anti-fat bias. *Health Psychology*, *22*, 68–78.

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

Zimmerman, D. W., & Williams, R. H. (2000). Restriction of ranger and correlation outlier-prone distributions. *Applied Psychological Measurement*, *24*, 267–280.

Author biographies

SARAH RODDY is currently completing her PhD at the School of Psychology, National University of Ireland Galway. Her doctoral research, which was funded by a scholarship from the Irish Research Council for the Humanities & Social Sciences, has involved comparing methodologies for the assessment of implicit cognition in the context of body shape-related attitudes.

IAN STEWART is a full time academic at the School of Psychology, National University of Ireland, Galway. His current research focuses on the analysis of language and cognition from a behaviour analytic and more specifically relational frame theory perspective. He has co-published a number of articles on the use of the Implicit Relational Assessment Procedure as a measure of implicit cognition.

DERMOT BARNES-HOLMES is foundation Professor of Psychology at the National University of Ireland, Maynooth and was Head of Department from 1999 to 2008. He has published widely in behavioural psychology with a particular focus on human language and cognition. He founded and currently leads the Implicit Cognition and Behaviour research group.