

UC Davis

UC Davis Previously Published Works

Title

Automatic and Controlled Components of Prejudice Toward Fat People: Evaluation Versus Stereotype Activation

Permalink

<https://escholarship.org/uc/item/1730r8rf>

Journal

Social Cognition, 18(4)

ISSN

0278-016X

Authors

Bessenoff, Gayle R
Sherman, Jeffrey W

Publication Date

2000-12-01

DOI

10.1521/soco.2000.18.4.329

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-NoDerivatives License, available at <https://creativecommons.org/licenses/by-nc-nd/4.0/>

Peer reviewed

AUTOMATIC AND CONTROLLED COMPONENTS OF PREJUDICE TOWARD FAT PEOPLE: EVALUATION VERSUS STEREOTYPE ACTIVATION

Gayle R. Bessenoff and Jeffrey W. Sherman
Northwestern University

This study investigated automatic and controlled components of anti-fat attitudes, the relationship between these components, and the extent to which each component predicts prejudicial behavior. Participants were primed with pictures of fat and thin women. Automatic activation of both evaluative responses and stereotypic knowledge were examined with lexical decision judgments on fat-stereotypical, thin-stereotypical, and stereotype-irrelevant trait words. Results showed greater automatic activation of negative evaluations to fat than thin women. Although, in general, automatic measures were found to be unrelated to self-reported anti-fat attitudes, one subcomponent of automatic evaluation was correlated with higher expressed dislike of fat persons. In addition, the automatic but not the controlled attitudinal measure predicted how far participants chose to sit from a fat woman. No stereotypicality effects were observed. Implications for reducing prejudice toward fat persons are discussed.

In recent years, considerable interest in research on stereotyping and prejudice has focused on the distinction between controlled and auto-

Gayle R. Bessenoff and Jeffrey W. Sherman, Department of Psychology, Northwestern University.

This research was supported by National Institute of Mental Health Grant 55037 to Jeffrey W. Sherman. We thank Josh Britton for his assistance in data collection. We would also like to thank Galen Bodenhausen, Wendi Gardner, Jim Sherman, Jeff Larson, Monica Biernat and three anonymous reviewers for their thoughtful comments on earlier versions of this article.

Correspondence concerning this article should be addressed to Gayle R. Bessenoff or Jeffrey W. Sherman, Department of Psychology, Northwestern University, 2029 Sheridan Road, Evanston, IL 60208-2710, e-mail: g-bessenoff@nwu.edu or sherm@nwu.edu

matic processes. Controlled, or conscious, processing occurs when the individual is aware of the process and has the time and the capacity to control or alter his or her responses. Automatic processes, on the other hand, occur without this knowledge and control, and require no conscious effort or attention (Shiffrin & Schneider, 1977). On some occasions, prejudicial behavior is based on explicit, consciously held, and controllable attitudes and stereotypes. At other times, however, prejudicial behavior may derive from implicitly held attitudes and stereotypes (e.g., associations between social groups and evaluations/stereotypes) that the perceiver may not be aware of, and that are applied unconsciously and without intent.

The present research was designed to further our understanding of the role of automatic and controlled processes in prejudicial attitudes and behavior. Our goal was to determine the relationships between automatic and controlled measures of both evaluations and stereotypes and how these factors independently and interactively affect behavior. We first examined the types of knowledge automatically activated in the presence of a member of a stigmatized and stereotyped group. Much information is associated with social groups, including cultural prototypes and stereotypes, as well as affective reactions. Does all this information automatically come to mind, or are only certain types of information strongly associated with the social group? Second, we investigated how these different types of automatically activated material correlate with explicitly expressed attitudes toward the target. Third, we examined the predictive quality of different types of automatically activated material as well as explicitly stated beliefs for prejudicial behavior.

AUTOMATIC ACTIVATION OF STEREOTYPES AND EVALUATIONS

Past research has demonstrated that stereotypical knowledge can be automatically activated, influencing subsequent judgments (e.g., Banaji & Greenwald, 1994; Banaji & Hardin, 1996; Banaji, Hardin, & Rothman, 1993; Blair & Banaji, 1996; Devine, 1989; Dovidio, Evans, & Tyler, 1986; Lepore & Brown, 1997). For example, Devine (1989) showed that an ambiguous target was judged to be relatively hostile when participants had been unconsciously exposed to a large number of words related to the cultural stereotype of African Americans. Inter-group attitudes have also been shown to be automatically activated in the presence of group members. Fazio, Jackson, Dunton, & Williams (1995) found that Caucasian participants made judgments about the valence of negative words following pictures of African American

faces faster than following Caucasian faces, whereas African American participants responded faster to negative words following pictures of Caucasians (see also, Dovidio, Kawakami, Johnson, Johnson, & Howard, 1997). This suggests automatic activation because participants were unaware of the relationship between the pictures and the subsequent task and did not have enough time to control their responses.

Most research on automatic and controlled components in prejudice has focused on either the activation of stereotype content (e.g., Banaji & Greenwald, 1995; Banaji & Hardin, 1996; Banaji et al., 1993; Blair & Banaji, 1996; Devine, 1989; Lepore & Brown, 1997) or evaluations (e.g., Dovidio et al., 1997; Fazio et al., 1995) in isolation. There may be, however, important interactions between stereotypes and evaluations. Wittenbrink, Judd, and Park (1997; see also Kawakami, Dion, & Dovidio, 1998) used a priming paradigm to examine the differential automatic activation of stereotypes and evaluations in response to a racial target. Participants were subliminally primed with the words "black" and "white," followed by word-nonword (lexical decision) judgments on positively and negatively valenced traits that were either stereotypical of African Americans or stereotypical of Caucasians. Results indicated that both stereotypes and evaluations were involved in automatic processes of prejudice. Thus, stereotypical and negative traits were both identified more quickly following the 'black' prime than the 'white' prime. However, Wittenbrink and colleagues (1997) also found an interaction between stereotypes and valence, such that there was greater facilitation for negative African American stereotypes following African American primes and positive Caucasian stereotypes following Caucasian primes. They suggested that different racial stereotypes might be associated with different valences, and termed the measure of these valenced stereotypes "implicit prejudice."

Locke, MacLeod, and Walker (1994) also investigated the interaction of automatic stereotypes and evaluations. They examined patterns of activation of positive and negative trait information either related or unrelated to the cultural stereotype of Australian Aborigines. These patterns of activation were examined while the participants were making judgments of the target group using a word-color naming interference task (Stroop). When the participants were familiar with the stereotypes of Aborigines, automatic activation of stereotype content was found. However, when participants were unfamiliar with these stereotypes, negative concepts were automatically activated, irrespective of their relationship to the stereotype.

THE RELATIONSHIP BETWEEN AUTOMATIC AND CONTROLLED MEASURES OF ATTITUDES

A second important issue to consider is how automatic activation relates to explicit forms of prejudice. Devine (1989) theorized that automatic and controlled processes in prejudice operate independently, such that automatically activated stereotypes are independent of explicit attitudes toward the target group. According to this model, high- and low-prejudiced individuals differ in their expressed beliefs toward stereotyped groups, but have equal knowledge of the cultural stereotypes, which are automatically activated in the presence of a member of the group. In support of this model, Devine (1989, Study 2) found that automatic stereotyping effects were equally strong for participants expressing both high and low prejudice on an explicit measure of prejudice.

In contrast, Fazio and others (1995; see also Dovidio et al., 1997; Hense, Penner, & Nelson, 1995; Kawakami et al., 1998; Wittenbrink et al., 1997) found significant individual differences in the automatic activation of prejudice. These data demonstrate that not everyone has the same race-related material automatically activated in memory. Moreover, Fazio et al. (1997) found that when motivation to appear nonprejudiced was low, automatic and controlled attitudinal responses were related. That is, for participants unconcerned with openly expressing racial attitudes, those high in prejudice had more negative automatic evaluations than those low in prejudice. No such correlation existed for participants uncomfortable expressing racial beliefs. Wittenbrink and colleagues (1997) also found a positive correlation between their measure of implicit prejudice and explicit beliefs. Although they did not measure motivation to control prejudice, under the assumption that people are generally motivated to not appear racially biased (e.g., Gaertner & Dovidio, 1986; McConahay, 1986), their data suggest that low motivation may not be necessary to observe correlations between measures of automatic and controlled prejudice.

PREDICTING BEHAVIOR

Of course, one of the central goals of measuring stereotypes and prejudice (or any attitude) is to be able to predict behavior. The distinction between relatively automatic and controlled prejudice may be of particular importance in this regard. Historically, explicitly held attitudes have not been very good predictors of prejudicial behavior (e.g., LaPiere, 1934; Wicker, 1969). However, for a variety of reasons, automatically activated knowledge may be expected to better predict behavior than explicitly reported attitudes. First, social norms against expressing

prejudice can contribute to people's reluctance to overtly admit negative attitudes or stereotypes. As such, explicit measures of prejudice may not always be accurate, and therefore would not predict behavior very well. In contrast, such self-presentational concerns will have little impact on measures that tap automatically activated knowledge. Second, people may not even be aware of the underlying attitudes and stereotypes they hold, and therefore could not accurately report them even if they wanted to do so (Wegner & Bargh, 1998). Measures of automatic activation avoid this difficulty. Finally, the influence of consciously held beliefs depends on perceivers having sufficient motivation and ability to enact them (e.g., Devine, 1989). If perceivers are either unmotivated or unable to enact their explicit beliefs, then automatically activated knowledge would be much more likely to influence behavior.

In fact, recent research has shown that some kinds of behaviors are better predicted by automatically activated attitudes than explicitly stated ones. For example, Fazio et al. (1995) found that participants with greater activation of negative evaluation toward African Americans also exhibited behaviors such as decreased smiling toward and greater spatial distance from an African American experimenter. Controlled measures of prejudice did not predict these behaviors. Similarly, Dovidio et al. (1997) found automatic evaluations, but not controlled attitudes, to be related to nonverbal visual measures of negative arousal such as blinking and eye contact toward an African American interaction partner. In contrast, Dovidio et al. (1997) found that controlled, but not automatic, attitudes predicted judgment ratings, such that participants who expressed more negative beliefs about African Americans rated the African American partner significantly more negatively than a Caucasian partner. Automatically-activated attitudes did not predict these ratings. Such differentiation of behavioral prediction was found even when automatic and controlled measures of attitudes were positively correlated (Dovidio et al., 1997, Experiment 2). This led Dovidio et al. (1997; see also the MODE model, Fazio, 1990) to propose that the processing level of attitudes (automatic or controlled) that better predicts behavior depends on the type of behavior. Specifically, they theorized that automatically activated attitudes better predict relatively spontaneous and unconscious behaviors (such as nonverbal behaviors), whereas self-reported attitudes better predict more deliberate responses (see also Devine, 1989).

EXPANDING ON PREVIOUS LITERATURE

As previously described, research into prejudicial attitudes and behavior primarily has looked at the activation and effects of either stereo-

types or evaluations independently. Although Wittenbrink et al. (1997) demonstrated that valenced stereotypes were automatically activated by racial group labels and correlated with explicit beliefs, valenced traits unrelated to the stereotype were not included in Wittenbrink et al.'s (1997) design. This limited their ability to examine the generality of the valence effect outside the stereotype. The current study addressed this issue by using a design that included not only positively and negatively valenced target traits related to each target group's stereotype, but also valenced traits not associated with the stereotype of the target group.

In addition, Wittenbrink et al. (1997) did not examine the effects of the automatic activation of valenced stereotypes on behavior. To more fully understand attitudinal processes, it is important to examine not only the activation of stereotypes and evaluations, but also their independent and interactive effects on behavior. It may be that, although both evaluations and stereotypes are automatically activated, they are differentially related to behavior. Therefore, in addition to measuring automatic and controlled attitudes, we also collected a behavioral measure in order to examine the relative influence of both stereotypes and evaluations on prejudicial behavior.

Thus, the current study had three goals: (1) to determine the specific informational content (stereotypes as well as evaluations) automatically activated in the presence of a stereotyped target; (2) to examine the relationship between this automatic activation and explicit beliefs about the target's group; and (3) to determine the role of both automatically activated and controlled knowledge in nonverbal prejudicial behavior toward a member of this stereotyped group.

In addition to addressing these basic issues, a further contribution of our research is that it examined these attitudinal processes with a different type of social group than previously examined. Most studies examining automatic and controlled intergroup perception have focused on racial and gender-based stereotypes and attitudes. However, the accurate assessment of beliefs about these groups is difficult to achieve. In particular, because the expression of negative attitudes toward racial and gender groups is socially undesirable, such expressions are likely to be moderated by self-presentational concerns (Dunton & Fazio, 1997; Plant & Devine, 1998). This may increase the difficulty of assessing the relationships among expressed beliefs, automatically activated knowledge, and behavior (Fazio et al., 1995). Other stigmatized groups exist that do not appear to generate such social desirability concerns.

One such group is fat people¹. Anti-fat attitudes have been found to be similar in nature to racial attitudes, in that they are negative and discriminatory (Allon, 1982; Yucker & Allison, 1994). Those who are fat are consistently described more negatively than those with other disabilities

(Richardson, 1993), receive less financial support for schooling from their parents (Crandall, 1991), and are discriminated against in the job market (Rothblum, Miller, & Garbutt, 1988). These sanctions are even more severe for women (Orbach, 1978; Oswalt & Davis, 1990; Rodin, Silberstein, & Striegel-Moore, 1984). In contrast to racial prejudice, however, negative attitudes about those who are fat are not associated with strong normative pressure to appear non-prejudiced (Crandall, 1994).

OVERVIEW AND PREDICTIONS

A modified version of a priming procedure similar to that used by Wittenbrink et al. (1997; see also Perdue, Dovidio, Gurtman, & Tyler, 1990) was used to assess what informational and evaluative content is automatically activated in the presence of fat women. Participants were presented with multiple picture–word pairs on a computer. The primes were pictures of fat women, thin women, and neutral items—and target words consisted of valenced traits. In order to assess stereotyping and evaluation effects separately, the target stimuli consisted of positive and negative traits that were fat-stereotypic, thin-stereotypic, and irrelevant to either stereotype. The extent of automatic activation was determined by response times on making word–nonword judgments on the trait words, depending on the nature of the pictures that preceded them. Based on previous research on automatic activation of attitudes toward other stigmatized groups (Dovidio et al., 1997; Fazio et al., 1995), as well as work suggesting the existence of strong anti-fat attitudes (e.g., Crandall, 1991; Richardson, 1993; Rothblum et al., 1988; Yucker & Allison, 1994), we expected that activation should be greater for negative trait words when preceded by a picture of a fat woman as compared to a thin woman. We also predicted that, collapsed across positive and negative valenced traits, activation should be greater for fat-stereotypical words following pictures of fat women as compared to thin women (e.g., Blair & Banaji, 1996; Devine, 1989; Dovidio et al., 1986). Finally, based on Wittenbrink et al.'s (1997) findings, we expected particularly high levels of activation for negative fat-stereotypical trait words following fat primes as compared to thin primes, as well as activation for positive thin-stereotypical trait words following thin as compared to fat primes.

The study also examined the relationship between automatically activated knowledge and explicit beliefs about fat people. To accomplish this, participants were preselected based on explicit anti-fat attitudes. Because motivation to control prejudice against fat people is relatively low (Crandall & Biernat, 1990), we expected to find a positive correlation between explicitly expressed attitudes and knowledge automatically activated in the priming task (Fazio et al., 1995).

However, despite this predicted relationship between automatic and controlled attitude measures, we expected them to differentially predict behavior. To measure behavior, we measured participants' preferred seating distance from a fat woman. Nonverbal measures of social distance are generally assumed to tap relatively unconscious behaviors (see Crosby et al., 1980; Dovidio et al., 1997; Fazio et al., 1995; Word, Zanna, & Cooper, 1974). Moreover, people are relatively unaware of the implications of this kind of behavior as far as prejudice is concerned. As such, seating behavior was expected to be better predicted by the automatic activation of prejudice than by explicit responses (Devine, 1989; Dovidio et al., 1997; Fazio, 1990). Therefore, in our study, participants with more negative automatic attitudes were expected to distance themselves farther away from the fat woman. It was unclear what role activated stereotype content would have in prejudicial behavior, as the automatic activation of stereotype content has not been directly tested in previous research as a predictor of prejudicial behavior. It might have a direct effect, or, following Wittenbrink et al. (1997), behavior may be influenced by an interaction between automatically activated evaluations and stereotypes.

METHOD

PARTICIPANTS

Participants were 127 undergraduates (64 female and 63 male) recruited from Northwestern University's Introductory Psychology participant pool, which offers course credit for participation. Participants were chosen based on their scores on two subscales of Crandall's (1994) 13-item Anti-fat Questionnaire (Dislike and Willpower), administered during group testing at the beginning of the term.² This questionnaire measures aspects of attitudes toward fat people (e.g., "I really don't like fat people much") on a 10-point Likert scale (0 = strongly disagree to 9 = strongly agree). Participants with the most extreme low scores (range = 0–2) on the scale were selected as low-prejudiced participants in the study. Due to the fact that there were very few individuals with extremely high scores, participants with scores 4 and above were selected as high-prejudiced participants (range = 4–6.6). However, it was later determined that the Dislike Subscale alone was a more appropriate measure of prejudice against fat persons, as Willpower is considered a separate construct (Crandall, 1994). Therefore, a median split based on scores on the Dislike Subscale was used to determine which participants were high-prejudiced and which were low-prejudiced (total range = 0–7.33, median = 1.83; for low prejudice, range = 0–1.83; $M = .71$, $SD = .51$; for high preju-

dice, range = 1.86–7.33, $M = 3.69$, $SD = 1.10$).³ The experimenter was blind to participants' prejudice level. Experimental sessions included 1 to 4 people for the priming procedure. Participants were tested individually for the behavioral measure.

Of these 127, one participant's data was excluded due to an abnormally high level of lexical decision errors (over 50%) on the priming task. Because 8 of the participants were assigned to the study after mass testing, they did not have explicit attitude data, and were not included in analyses involving this variable. The data from 12 participants were not included in the correlational analyses with behavior; 8 participants did not complete this measure properly (either due to participant or experimenter error), and the data of 4 participants were excluded due to admission of doubting the cover story about interacting with another student.

STIMULUS MATERIALS

PICTURE PRIMES

For the photo primes, photographs of 6 different 20-something thin women were cut out from magazines and scanned into the computer. The pictures of the thin women were manipulated using a computer liquid-imaging program [Kai's Power Goo, 1996] to create fat versions of the same women. Therefore, there were 12 pictures of women in all: 6 thin and 6 fat. The pictures were pretested for body type on an 8-point scale (1 = extremely underweight to 9 = extremely overweight). Fat pictures were seen as significantly more fat (mean rating = 7.26) than the thin pictures (perceived around average; mean rating = 3.83), $F(1, 22) = 351.46$, $p < .001$. There were also 6 pictures of neutral objects (mug, chair, sunglasses, camera, baseball cap, lamp) used as baseline primes.

STIMULUS WORDS

Positive and negative fat-stereotypic, thin-stereotypic, and weight stereotype-irrelevant words were assembled from traits generated by pilot participants, as well as from Butler, Ryckman, Thornton, and Bouchard's (1993) list of trait words highly associated with endomorph (fat) and mesomorph (average) body types. These trait words were then rated by 48 pilot participants in terms of how well they characterized either "fat women" or "slim women" on a 6-point scale (1 = not at all to 6 = very). On the basis of these ratings, a critical set of 36 trait words was selected for use in the experiment. These include 6 negative fat-stereotypic traits (e.g., insecure), 6 positive fat-stereotypic traits (e.g., kind), 6 nega-

tive thin-stereotypic traits (e.g., selfish), 6 positive thin-stereotypic traits (e.g., confident), 6 negative stereotype-irrelevant traits (e.g., violent) and 6 positive stereotype-irrelevant traits (e.g., musical). Trait words that were rated high for fat (mean rating = 4.12 for positive, 3.92 for negative), but not thin (mean rating = 3.38 for positive, 2.67 for negative) were used for the fat-stereotypic words. Trait words that were rated high for thin (mean rating = 4.42 for positive, 3.88 for negative) but not fat (mean rating = 2.73 and 2.55 for negative) were used for the thin-stereotypic words. Stereotype-irrelevant words were chosen from traits that were rated as pertaining equally to fat (mean ratings = 3.48 for positive, 2.72 for negative) and thin women (mean ratings = 3.67 for positive, 3.10 for negative). The resulting 6 subsets of traits are presented in the Appendix. Eighteen nonwords were also used for filler trials, created by mixing up the letters of the other words (e.g., *gerenetic*, *triactive*).

APPARATUS

The experimental task was administered on Macintosh Performa 475 microcomputers running SuperLab software (1994). All stimuli were presented in the center of the computer screen. All pictures were approximately 3 cm in height, slightly degraded, and presented in black and white on a white background. Words were black on a white background, in Geneva size 14 font.

PRIMING PROCEDURE

The picture primes were presented for 15 ms. Each prime was presented with a backward and forward mask, consisting of a conglomerate of pieces taken from each picture. Although participants indicated that they were sometimes aware that pictures had been presented, the picture primes could only be identified as human figures; it was very difficult to tell whether the person in the picture was fat or thin in the 15 ms interval. There was a delay of either 450 or 1000 ms before the presentation of the test item for 250 ms, which consisted of either a nonword or a word from one of the trait categories described above.⁴ Participants made a lexical decision on the test item, pressing one key if the test item was a real word and another if it was a nonword. The dependent measure consisted of the response latencies in making these judgments.

The picture-word pairs were presented randomly and were divided into 6 blocks of trials, each consisting of 27 pairs, with a short break in between. Picture-word pairs were counterbalanced such that each of the 54 words was paired with a picture from each of the 3 categories (fat, thin, object). Both versions (fat and thin) of each woman were paired

with the same words, consisting of at least 2 trait words from each of the six categories. Two counterbalancing versions were used, varying which words were assigned to which individual female picture.

PROCEDURE

Participants worked in individual rooms; all instructions were given verbally by the experimenter. The experiment was presented as an investigation into verbal and communication skills; no mention was made of the role of fat stereotypes or attitudes. For the first task, participants were informed that their verbal skills were being assessed. This task consisted of the priming procedure outlined above. Participants were instructed that they would see an orienting stimulus flash quickly on the screen, followed by a string of letters. The orienting stimulus was the presentation of the picture prime. Participants were instructed to make a word-nonword judgment on the letter string as quickly and as accurately as possible. The "F" and "J" keys on the keyboard were used for this purpose, and were labeled "no" and "yes," respectively. The time it took to respond to the stimulus was recorded, as well as the accuracy of the judgment.

After completing the priming task, participants completed a puzzle filler task for 5 minutes. They were then given a folder containing information describing another student they were to meet. This folder included a photograph of a 200+ pound 20-year-old woman and rudimentary information, such as her name, ID number, and major. Participants were told that they would be taken into another room to meet and converse with this student in order to assess the application of verbal skills in a real-time interaction. The other room contained a solitary chair on which a coat and backpack had been placed. Participants were informed that the other student had been sitting in this chair, but had momentarily stepped out of the room. Participants were then asked to bring in another chair and sit and wait until the other student returned. The dependent measure of behavior consisted of the distance the participant chose to place his or her chair from the other "student's" chair (see Macrae, Bodenhausen, Milne, & Jetten, 1994). Greater seating distances indicated greater prejudicial behavior.

RESULTS

Presented first is the analysis of the latencies from the priming task to examine automatic activation following fat and thin female primes. Subsequently, we discuss relationships between the automatic responses and

the questionnaire measures. Finally, behavioral predictors are examined.

ANALYSIS OF RESPONSE LATENCIES

For each participant, response times for each test word preceded by the fat, thin, and neutral picture primes were recorded. Trials on which participants incorrectly identified the test word and latencies greater than 1 sec were removed from the analysis (5% of the responses; see Ratcliff, 1993). Information automatically activated by fat and thin women was determined by comparing response latencies (RTs) to items following thin and fat picture primes to those following neutral primes. RTs for each word to the fat prime and the thin prime were subtracted from those of the neutral baseline. Therefore, larger numbers indicate facilitation from baseline. Words from each of the 6 trait categories (positive fat-stereotypic, negative fat-stereotypic, positive thin-stereotypic, negative thin-stereotypic, positive stereotype-irrelevant, negative stereotype-irrelevant) were averaged separately for both the fat and thin primes, creating 12 separate variables for each participant.

We conducted a 2 (Prime: fat, thin) \times 3 (Trait Stereotypicality: fat-stereotypic, thin-stereotypic, stereotype-irrelevant) \times 2 (Trait Valence: positive, negative) \times 2 (Prejudice Level: high, low) \times 2 (SOA: 1015, 465) \times 2 (Participant Gender: male, female) mixed-factorial ANOVA. This analysis yielded only two reliable effects including the between-subjects factors: a Prime \times Gender interaction, $F(1, 110) = 5.88, p < .05$, such that men and women responded differently to trait words preceded by fat primes than thin primes (men: $M_s = 4.11$ ms for fat primes, -5.12 ms for thin primes; women: ($M_s = 1.36$ ms for fat primes, 3.62 ms for thin primes for women), and a Trait Valence \times Prejudice Level, $F(1,110) = 4.11, p < .05$, such that high- and low-prejudiced participants responded differently to positive trait words than to negative trait words (regardless of prime; high: $M_s = -1.75$ ms for positive traits, 1.94 ms for negative traits; low: $M_s = 8.36$ for positive traits, -4.58 for negative traits).

EVIDENCE FOR AUTOMATIC EVALUATION

The analysis yielded only one theoretically significant effect, a Prime \times Trait Valence interaction, $F(1,110) = 4.02, p < .05$. Further analyses showed that greater facilitation occurred when negatively valenced traits had been preceded by fat primes ($M = 2.93$ ms) than thin primes ($M = -5.57$ ms), $F(1, 110) = 5.63, p = .02$, but not when positive traits had been preceded by the thin than fat primes ($M_s = 4.07$ and 2.53 , respectively), $F(1, 110) = .22, ns$ (see Figure 1). The ANOVA revealed no effects of trait

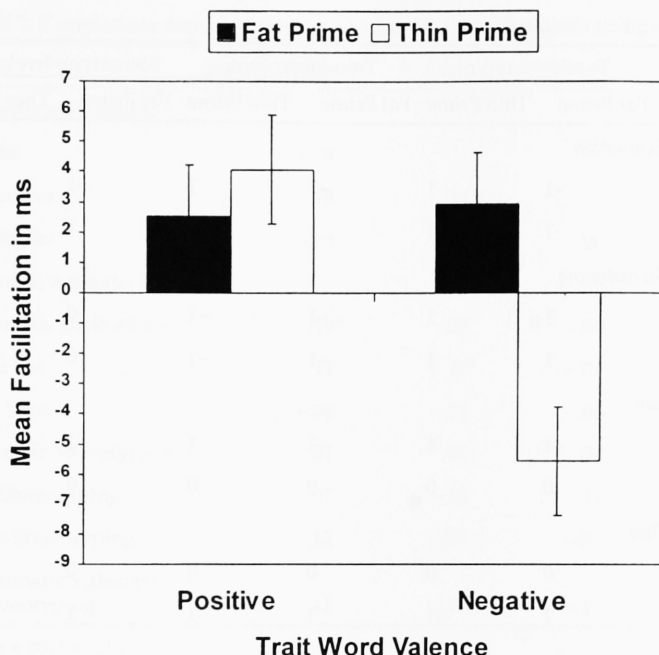


FIGURE 1. Mean response facilitation (in milliseconds) for positive and negative trait words preceded by fat and thin female picture primes. Error bars indicate standard error from the mean.

stereotypicality (all F s < 1.9).⁵ These data demonstrate that, in general, automatic attitudes toward fat women are significantly more negative than those toward thin women.

IMPLICIT CONTRASTS

In addition to the above tests, we conducted a series of contrasts to pinpoint the source of the automatic evaluation effect. These contrasts are outlined in Table 1. For completeness, the Automatic Evaluation Contrast is depicted (this is identical to the Prime X Target Valence interaction from the ANOVA results reported earlier), as is the Automatic Stereotyping contrast (not significant, as indicated in the ANOVA). A Positivity Bias Contrast examined whether responses to positive items differed for fat and thin primes, and was not found to be significant, $F(1,110) = .22$, ns. The Negativity Bias contrast, however, was found to be reliable, $F(1,110) = 5.63$, $p < .05$, such that participants responded significantly faster to negative items following fat than thin primes. An

TABLE 1. Implicit Contrasts

Measure	Fat-Stereotypic		Thin-Stereotypic		Stereotype-Irrelevant	
	Fat Prime	Thin Prime	Fat Prime	Thin Prime	Fat Prime	Thin Prime
<i>Automatic Evaluation</i>						
Positive	-1	1	-1	1	-1	1
Negative	1	-1	1	-1	1	-1
<i>Automatic Stereotyping</i>						
Positive	1	1	-1	-1	0	0
Negative	1	1	-1	-1	0	0
<i>Positivity Bias</i>						
Positive	-1	1	-1	1	-1	1
Negative	0	0	0	0	0	0
<i>Negativity Bias</i>						
Positive	0	0	0	0	0	0
Negative	1	-1	1	-1	1	-1
<i>Anti-Fat</i>						
Positive	1	0	1	0	1	0
Negative	-1	0	-1	0	-1	0
<i>Pro-Thin</i>						
Positive	0	1	0	1	0	1
Negative	0	-1	0	-1	0	-1
<i>Automatic Valenced Stereotyping</i>						
Positive	-1	0	0	1	0	0
Negative	1	0	0	-1	0	0

Anti-Fat contrast examining the facilitation difference due to valence only following the fat primes was not found to be reliable, $F(1, 110) = .007$, ns. The Pro-thin contrast, however, which examined the facilitation difference due to valence following the thin primes, was found to be reliable, $F(1, 110) = 4.26$, $p < .05$, such that participants responded significantly faster to positive than negative items following thin primes (compared to baseline).

Because Wittenbrink et al. (1997) found a Prime \times Valence \times Stereotypicality interaction without the inclusion of stereotype-irrelevant traits, we defined a seventh contrast—Automatic Valenced Stereo-

TABLE 2. Correlations Among Automatic, Controlled, And Behavioral Measures

Measure	Seating Distance	Dislike	Willpower	Fear of Fat
Self-Reported Attitude				
Dislike	-.08			
Willpower	-.19	.34**		
Fear of Fat	-.03	.23*	.08	
Automatic Attitude				
Automatic Evaluation	.19*	.09	-.09	.09
Anti-Fat	.11	.19*	-.04	.02
Pro-Thin	-.04	-.13	-.05	.08
Automatic Stereotyping	.01	.04	.07	.03
Fat Stereotyping	-.07	-.10	.11	-.01
Thin Stereotyping	.12	.07	-.03	-.16
Automatic Valenced Stereotyping	.11	.10	-.13	.04

Note. * $p < .05$, ** $p < .01$

typing—to examine this effect in our data. This contrast tests the specific valenced stereotypic associations for each group prime, replicating Wittenbrink et al.'s (1997) Implicit Prejudice measure, examining whether facilitation is stronger for negative fat-stereotypical traits following fat primes and positive thin-stereotypical traits following thin primes. This contrast also did not yield reliable results, $F(1, 110) = 1.67$, ns.

RELATIONSHIPS BETWEEN AUTOMATIC AND CONTROLLED RESPONSES

In the ANOVA described above, participants high and low on explicit prejudice did not differ in their automatic associations to the picture primes. To further examine the relationships between explicit measures and the automatic activation of stereotypes and evaluations, we correlated all three subscales of Crandall's (1994) Anti-Fat Questionnaire (Dislike, Willpower, Fear of Fat) with various patterns of automatic activation implied by the contrasts outlined in Table 1 (see Table 2).⁶ None of the automatic measures correlated with any of the subscales of the questionnaire, with the exception of the Anti-Fat component of Automatic Prejudice, which was found to be reliably related to Explicit Prejudice,

$r = .19, p < .05$. Individuals overtly expressing dislike for fat people also exhibited greater activation of negative than positive traits to fat primes.

BEHAVIOR

Prejudicial behavior was determined by measuring the distance participants placed their chair from the chair in which they believed the fat target would be sitting. Behavioral scores were recorded as the number of inches between the fronts of the two chairs. Correlational analyses were computed to examine self-reported anti-fat attitudes as well as automatically activated evaluations, stereotypes, and valenced stereotypes as predictors of this seating behavior (see Table 2).

Our measure of explicit anti-fat attitudes was not found to be correlated with behavior. In fact, none of the three separate subscales of Crandall's (1994) Anti-fat Questionnaire (Willpower, Dislike, Fear of Fat) were found to be significantly correlated with behavior (although the Willpower Subscale was marginally negatively correlated with seating distance, such that participants who endorsed the ideology that weight is controllable chose to sit closer to a fat woman).

Automatic Evaluation was found to be correlated with behavior, such that participants with greater negative than positive activation for the fat primes compared to the thin primes also chose to sit farther away from a fat woman. The Anti-Fat and Pro-Thin components alone were not reliably related to behavior, demonstrating that such behavior toward a fat woman is based not on automatic evaluations toward fat or thin people, but on differential automatic evaluations toward the two groups. Neither Automatic Stereotyping nor Automatic Valenced Stereotyping was found to be a reliable predictor of behavior. Seating behavior was also not related to activation of either the Fat- or Thin-Stereotyping measure. Therefore, spontaneous social distance behavior was predicted by automatic negative evaluation, but not by measures of automaticity that included stereotype activation, nor by expressed beliefs. Overall, these results suggest that automatic anti-fat attitudes exist outside the boundaries of the fat stereotype, and that unconscious, nondeliberative prejudicial behavior is best predicted by these automatic attitudes.

DISCUSSION

Data from this research suggest that automatically activated attitudes toward fat women are significantly more negative than those for thin women. There was a main effect of valence such that there was greater activation for negative traits upon exposure to fat than thin female

primes. That this pattern was not qualified by an interaction with stereotyping demonstrates that differences in negative evaluation associated with fat and thin women is relatively diffuse, and extends beyond fat-relevant traits.⁷

Based on Crandall's (1994) suggestion that there are few social desirability norms about expressing negative attitudes towards fat people, we expected that automatic and controlled anti-fat attitudes might be related in our study (e.g., Fazio et al., 1995). This was true to a limited extent. Although, in general, the automatic measures were found not to be related to self-reported anti-fat attitudes, the Anti-Fat sub-component of Automatic Evaluation did significantly correlate with Explicit Prejudice, such that participants with higher expressed dislike of fat people also had greater automatic activation of negative than positive traits to fat primes (as compared to a neutral baseline). We do not know why only the Anti-Fat component, but not the overall or Pro-Thin component was correlated in this way. One possibility may be that the Anti-Fat Questionnaire reflects attitudes toward fat persons without reference to thin persons. As such, automatic measures that include reference to thin persons would not be correlated with explicit measures.

Overall, the general lack of a relationship between automatic and controlled processes across our different measures suggests that, even when motivation to control prejudice is assumed to be low, correlations between automatic and controlled attitudes may be weak. There is a good reason why this may occur. Namely, because people may not be able to accurately report their underlying attitudes and stereotypes (Wegner & Bargh, 1998), measures of automatic and controlled attitudes may remain uncorrelated, even when participants are not motivated to appear nonprejudiced. However, because we did not directly measure motivation in our study, we refrain from making strong claims about the implications of our data for clarifying the relationship between motivations and automatic/controlled correlations.

Research on the extent to which automatic and controlled attitudes and stereotypes are related has been mixed, and consistent moderators of the relationship have been difficult to identify (e.g., Banaji & Greenwald, 1995; Devine, 1989; Dovidio et al., 1997; Fazio et al., 1995; Wittenbrink et al., 1997). Our data suggest that this may be due, in part, to the wide variety of measures used to assess automatic and controlled components of prejudice. Though almost all of our measures showed no relationship, a few subcomponents were correlated. Thus, conclusions about the relationship between automatic and controlled aspects of attitudes and stereotypes may be dependent on the measures used. Obviously, this will be an important issue for future research.

Our measures of automatic and controlled components of anti-fat atti-

tudes were also found to differentially predict behavior. Automatic, but not controlled, responses were found to be positively correlated with unconscious nonverbal behavior, such that participants with greater activation of negative than positive traits to fat than thin female primes chose to sit farther from the fat woman. This is consistent with Fazio's (1990; Fazio et al., 1995) MODE model, which suggests that unconscious aspects of an attitude best predict spontaneous forms of behavior, whereas conscious aspects best predict deliberative responses (see also Dovidio et al., 1997). To the extent that nonverbal behaviors such as physical proximity are good examples of unconscious behaviors, and are not based on deliberative processing of one's attitude toward the object (see Dovidio, Brigham, Johnson, & Gaertner, 1996), they should be singly influenced by automatically activated evaluations (Fazio & Dunton, 1997). Certainly, the act of placing a chair is a controlled action. However, in order to control that action in terms of prejudicial behavior, one must be aware not only of the behavior, but that it can be construed as something in need of regulation (e.g., a prejudicial response). We do not believe that participants in this study recognized this behavior as such. Therefore, we construed seating distance as unconscious—or spontaneous—prejudicial behavior. Nonetheless, an important follow-up to the present study should look at clearly deliberative as well as unconscious behaviors toward fat people.

Interestingly, our attitude-behavior correlation was not only based on automatic evaluative responses, but on the relative activation of evaluations of fat women as compared to thin women. Absolute automatic negative evaluation toward fat women (relative to a neutral baseline), although correlated with expressed dislike for fat persons, did not predict spontaneous prejudicial behavior toward a fat woman, nor did automatic positive evaluation toward thin women.

In contrast to past findings, no effects of item stereotypicality were observed in this research (but see Note 5). We also found no evidence of the effect of automatic activation of valenced stereotypes demonstrated by Wittenbrink et al. (1997). One possible reason for this discrepancy is that past research on automatic stereotyping has primarily used word primes (e.g., Banaji & Hardin, 1996; Banaji et al., 1993; Blair & Banaji, 1996; Devine, 1989; Dovidio et al., 1986; Wittenbrink et al., 1997), whereas the present study used pictures of group members. It may be that lexical stereotype labels are particularly strong cues for the activation of stereotyped traits (e.g., Gilbert & Hixon, 1991; Lepore & Brown, 1997). By contrast, pictorial stimuli may be more likely to induce automatic affect than words, which may then drive the evaluative response. Other possible explanations for this discrepancy between the current study and past work have to do with the fact that the stereotypes used in

the present research are different in many ways (e.g., content, strength, social desirability) from those previously examined (which have been mostly racial stereotypes). It may be that stereotypes of the overweight are simply weaker, and therefore less associated with the target group than is true for racial stereotypes.

Individual differences in stereotype activation also did not predict behavior. This may suggest that stereotypes play little role in determining spontaneous, nonverbal behavior. However, based on this one set of data, such a conclusion would seem to be premature. It may be that different types of behavior, or their relative contexts, are influenced by different types of automatically activated information. Our measure of social distance was in the domain of a neutral interaction. By contrast, if being fat was relevant to the situation (e.g., appearance related), activated stereotypes may have accounted for some portion of the behavior. In addition, the automatic activation of other kinds of stereotypes (e.g., racial) may be better predictors of behavior. Clearly, more research is needed to examine the relationship between stereotype activation and behavior.

IMPLICATIONS AND CONCLUSION

Researchers' attempts to reduce negative anti-fat attitudes have often focused on providing information to counter beliefs that fat people are lazy or lacking in willpower. Yet such an approach has been found to be less than ideal (e.g., Nichols, Waters, Woolaway, & Hamilton-Smith, 1988; Yucker & Allison, 1994). In part, this may be because automatically activated responses play a larger role than beliefs in determining people's unconscious behavioral reactions toward fat people. Moreover, it appears that the automatically activated information that best predicts such behavior may be evaluatively based. Such a relationship between automatic attitudes and behavior has great implications for discrimination, for it is the reactions over which we have little to no control or awareness that will produce discriminatory behavior in nondeliberative situations, whether or not we are motivated to act in a nonprejudiced way. Data from the current study demonstrate that automatic reactions to fat people are more negative than those toward thin people, and that the extent of this negativity is related to unconscious nonverbal behavior toward a fat woman. It appears, then, that attempts to lessen prejudice and discrimination may be profitably focused not only on changing negative stereotypes, but also on inhibiting the automatic activation of negative evaluation that occurs in the presence of members of stigmatized social groups.

APPENDIX. TRAIT WORDS USED IN LEXICAL-DECISION TASK

Positive Fat-Stereotypic

caring
friendly
humorous
kind
maternal
sympathetic

Negative Fat-Stereotypic

insecure
introverted
lazy
passive
unhealthy
unpopular

Positive Thin-Stereotypic

athletic
attractive
confident
disciplined
energetic
happy

Negative Thin-Stereotypic

aggressive
competitive
conceited
demanding
selfish
vain

Positive Stereotype-Irrelevant

artistic
clean
economical
hardworking
musical
orderly

Negative Stereotype-Irrelevant

boring
forgetful
greedy
jealous
rude
violent

NOTES

1. The words "fat" and "anti-fat" were chosen to be consistent with the terms used by Crandall (1994). No derogatory meaning is implied.

2. Participants were preselected based on scores in the top and bottom range of anti-fat prejudice. However, due to circumstances with the institutional review board, the questionnaire administered to the second group of participants was altered slightly from Crandall's (1994) original set of 13 questions. A few items were removed, and some questions with a more positive spin were added, taken from the Attitudes towards Obese Persons Scale (Allison, Basile, & Yucker, 1991), the Beliefs about Obese Persons Scale (Allison, Basile, & Yucker, 1991), and the Attitudes

Toward Obese Adult Patients Scale (Bagley, Conklin, Isherwood, Pechiulis, & Watson, 1989). However, only items consistent across both sets of questionnaires from Crandall's (1994) items were used in the analyses.

3. By basing explicit prejudice level on the Dislike Subscale, our sample is somewhat less skewed in relation to the population we sampled from. Our sample included 39.34% from the bottom third of the population, 13.11% from the middle third, and 46.72% from the top third.

4. The two SOA latency conditions (465 ms vs. 1015 ms) were conducted in consecutive school terms. Although acknowledgment is made to possible problems with the lack of random assignment between the two conditions, statistical analyses demonstrated no critical between-study differences. The SOA was varied in this way to examine potential differences in activation associated with participants' ability to engage in controlled processing following the brief, but supraliminal, presentation of the prime. Short SOAs do not permit the influence of conscious intentions on subsequent target responses. By contrast, SOAs over 1000 ms permit such controlled processing (e.g., Neely, 1977). The fact that no effects of SOA were found suggests that providing greater opportunity to control responses did not alter results. This is not terribly surprising, given that the pictures were very difficult to identify and that participants were not aware of the purpose of the study or that their lexical decision response times could be used to determine prejudice.

5. To further investigate automatic stereotype activation, we also conducted separate analyses on the fat-stereotypic, thin-stereotypic, and stereotype-irrelevant traits. Within the fat-stereotypical traits, there was a marginal main effect of Prime, $F(1,110) = 2.99, p = .087$, such that facilitation was greater to fat-stereotypical traits when preceded by the fat ($M = .40$ ms) than thin ($M = -8.30$ ms) picture primes. No significant effects were found in the separate analyses of the thin-stereotypical or stereotype-irrelevant traits.

6. Because we preselected for individuals with extreme scores, correlations involving the explicit measure may not generalize to individuals with scores in the mid-range of anti-fat prejudice.

7. The fat primes in this study were computer generated from pictures of thin women. Although care was taken to make the fat versions look as realistic as possible, a pre-test showed that the fat primes did appear less realistic and more altered (e.g., they seemed "stretched") than the original thin pictures (mean ratings = 4.17 & 2.71, respectively, on a scale from 1 = completely realistic to 6 = completely altered), $F(1, 22) = 10.13, p < .01$. It is conceivable that these differences could have contributed to the evaluative activation results found in this study if perceivers react more negatively to altered than nonaltered pictures. However, the primes

were presented for such a short amount of time (15 ms) that it was difficult to discern much more than that the silhouettes of the bodies were either fat or thin body shapes, if even that. The fact that participants judged the fat pictures as more altered when explicitly asked to make such judgments does not mean that the pictures were spontaneously perceived that way during the 15 ms presentation. Moreover, automatic evaluations were found to be related to behavior. It is not clear why negative reactions to pictures based on perceptions of alteration would predict behavior toward a fat woman. Thus, we believe it is highly unlikely that this factor can account for the data.

REFERENCES

- Allison, D. B., Basile, V. C., & Yucker, H. E. (1991). The measurement of attitudes toward and beliefs about obese persons. *International Journal of Eating Disorders, 10*, 599-807.
- Allon, N. (1982). The stigma of overweight in everyday life. In B. Wolman (Ed.), *Psychological aspects of obesity: A handbook* (pp.130-174). New York: Van Nostrand Reinhold.
- Bagley, C. R., Conklin, D. N., Isherwood, R. T., Pechliulis, D. R., & Watson, L. A. (1989). Attitudes of nurses towards obesity and obese patients. *Perceptual and Motor Skills, 68*, 954.
- Banaji, M. R., & Greenwald, A. G. (1995). Implicit stereotyping and prejudice. In M. Zanna & J. Olson (Eds.), *The psychology of prejudice: The Ontario symposium* (Vol. 7, pp. 55-76). Hillsdale, NJ: Erlbaum.
- Banaji, M. R., & Hardin, C. D. (1996). Automatic stereotyping. *Psychological Science, 7*, 136-141.
- Banaji, M. R., Hardin, C., & Rothman, A. J. (1993). Implicit stereotyping in person judgment. *Journal of Personality & Social Psychology, 65*, 272-281.
- Blair, I. V., & Banaji, M. R. (1996). Automatic and controlled processes in stereotype priming. *Journal of Personality and Social Psychology, 70*, 1142-1163.
- Butler, J. C., Ryckman, R. M., Thornton, B., & Bouchard, R. L. (1993). Assessment of the full content of physique stereotypes with a free-response format. *Journal of Social Psychology, 133*, 147-162.
- Crandall, C. S. (1991). Do heavy-weight students have more difficulty paying for college? *Personality and Social Psychology Bulletin, 17*, 606-611.
- Crandall, C. S. (1994). Prejudice against fat people: Ideology and self-interest. *Journal of Personality and Social Psychology, 66*, 882-894.
- Crandall, C. S. & Biernat, M. R. (1990). The ideology of anti-fat attitudes. *Journal of Applied Social Psychology, 20*, 227-243.

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

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

Dovidio, J. F., Brigham, J. C., Johnson, B. T., & Gaertner, S. L. (1996). Stereotyping, prejudice, and discrimination: Another look. In N. Macrae, M. Hewstone, & C. Stangor (Eds.), *Foundations of stereotypes and stereotyping* (pp. 276–319). New York: Guilford.

Dovidio, J. F., Evans, N. E., & Tyler, R. B. (1986). Racial stereotypes: The contents of their cognitive representations. *Journal of Experimental Social Psychology*, *22*, 22–37.

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

Dunton, B. C., & Fazio, R. H. (1997). An individual difference measure of motivation to control prejudiced reactions. *Personality and Social Psychology Bulletin*, *23*, 316–326.

Fazio, R. H. (1990). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 23, pp. 75–109). Orlando, FL: Academic Press.

Fazio, R. H., & Dunton, B. C. (1997). Categorization by race: The impact of automatic and controlled components of racial prejudice. *Journal of Experimental Social Psychology*, *33*, 451–470.

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

Gaertner, S. L., & Dovidio, J. F. (1986). The aversive form of racism. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 61–89). Orlando, FL: Academic Press.

Gilbert, D. T. & Hixon, J. G. (1991). The trouble of thinking: Activation and application of stereotypic beliefs. *Journal of Personality and Social Psychology*, *60*, 509–517.

Hense, R., Penner, L., & Nelson, D. (1995). Implicit memory for age stereotypes. *Social Cognition*, *13*, 399–415.

Kai's Power Goo [Computer software]. (1996). Carpinteria, CA: MetaTools, Inc.

Kawakami, K., Dion, K. L., & Dovidio, J. F. (1998). Racial prejudice and stereotype activation. *Personality and Social Psychology Bulletin*, *24*, 407–416.

LaPiere, R. T. (1934). Attitudes vs. actions. *Social Forces*, 13, 230–237.

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

Locke, V., MacLeod, C., & Walker, I. (1994). Automatic and controlled activation of stereotypes: Individual differences associated with prejudice. *British Journal of Social Psychology*, 33, 29–46.

Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Jetten, J. (1994). Out of mind but back in sight: Stereotypes on the rebound. *Journal of Personality and Social Psychology*, 67, 808–817.

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.

Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology*, 106, 226–254.

Nichols, S., Waters, W. E., Woolaway, M., & Hamilton-Smith, M. B. (1988). Evaluation of the effectiveness of a nutritional health education leaflet in changing public knowledge and attitudes about eating and health. *Journal of Human Nutrition & Dietetics*, 1, 233–238.

Orbach, S. (1978). *Fat is a Feminist Issue*. New York: Paddington Press.

Oswalt, R., & Davis, J. (1990). Societal influences on a thinner body size in children. *Proceedings and Abstracts of the Annual Meeting of the Eastern Psychological Association*. Philadelphia, PA.

Perdue, C. W., Dovidio, J. F., Gurtman, M. B., & Tyler, R. B. (1990). Us and them: Social categorization and the process of intergroup bias. *Journal of Personality and Social Psychology*, 59, 475–486.

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

Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin*, 114, 510–532.

Richardson, S. A. (1993). Children's values in regard to disabilities: a reply to Yuker. *Rehabilitation Psychology*, 28, 131–140.

Rodin, J., Silberstein, L. R., & Striegel-Moore, R. H., 1984. Women and weight: A normative discontent. In T. B. Sonderegger (ed.), *Psychology and gender: Nebraska Symposium on Gender, 1984* (pp. 267–307). Lincoln: University of Nebraska Press.

Rothblum, E. D., Miller, C. T., & Garbutt, E. D. (1988). Stereotypes of obese female job applicants. *International Journal of Eating Disorders*, 7, 277–283.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic hu-

man information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, 84, 127–190.

SuperLab [Computer software]. (1994). Phoenix, AZ: Cedrus Corp.

Wegner, D. M., & Bargh, J. A. (1998). Control and automaticity in social life. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology*, (4th ed., Vol. 2, pp. 446–496). New York: McGraw Hill.

Wicker, A. W. (1969). Attitudes versus actions: The relationship of verbal and overt behavioral responses to attitude objects. *Journal of Social Issues*, 25, 41–78.

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.

Word, C. O., Zanna, M. P., & Cooper, J. (1974). The nonverbal mediation of self-fulfilling prophecies interracial interaction. *Journal of Experimental Social Psychology*, 10, 109–120.

Yuker, H. E., & Allison, D. B. (1994). Obesity: Sociocultural perspectives. In L. Alexander-Mott & B. D. Lumsden (Eds.), *Understanding eating disorders* (pp. 243–270). Washington, DC: Taylor & Francis.