

The Linguistic Intergroup Bias As an Implicit Indicator of Prejudice

William von Hippel, Denise Sekaquaptewa, and Patrick Vargas

Ohio State University

Received: August 28, 1996; revised: January 2, 1997; accepted: January 8, 1997

The Linguistic Intergroup Bias (LIB) is the tendency to describe stereotypic events in more abstract terms than counterstereotypic events. We examined whether a paper-and-pencil measure based on the LIB could be used as an implicit indicator of prejudice. Experiment 1 demonstrated that a measure of implicit racial prejudice based on the LIB predicted whether subjects evaluated an African-American or a Caucasian as threatening. Experiment 2 extended this finding with indirect measures of threat. In both of these experiments, an explicit measure of prejudice failed to predict subjects' evaluations. Experiment 3 demonstrated that a measure of implicit gender prejudice based on the LIB was correlated with an implicit prejudice measure based on biased attributional processing. Across these three experiments, implicit and explicit measures of prejudice were largely uncorrelated. © 1997 Academic Press

The assessment of prejudice is plagued by problems associated with whether people are willing and able to report their negative attitudes toward certain groups. Indeed, most modern theories of prejudice are built around the notion that people are unwilling to admit, even to themselves, that they might hold negative attitudes toward members of certain groups (for a review, see Hilton & von Hippel, 1996). Thus, rather than directly asking people whether they like or dislike members of certain groups, explicit prejudice scales typically attempt to measure the theoretical correlates, causes, and consequences of prejudicial attitudes.

Despite their efforts to get around the "willing and able" problem, such self-report measures of prejudice are largely transparent to most people (e.g., Fazio, Jackson, Dunton, & Williams, 1995; Dunton & Fazio, 1997) and may be inherently limited in how far afield they can go while still maintaining construct validity (e.g., Sniderman & Tetlock, 1986). Thus, research has turned to the

We thank Marilyn Brewer, John Cacioppo, Steve Fein, Reid Hastie, Jon Krosnick, and Richard Petty for their helpful comments on earlier drafts of the manuscript.

Address correspondence and reprint requests to William von Hippel, Department of Psychology, Ohio State University, 1885 Neil Avenue, Columbus, OH 43210. E-mail: vonhippel.1@osu.edu.

implicit assessment of prejudice by examining behavioral responses (Dovidio, Brigham, Johnson, & Gaertner, 1996) and biases in judgment or decision time that are associated with prejudice (Banaji & Greenwald, 1995; Banaji & Hardin, 1996; Fazio et al., 1995; Wittenbrink, Judd, & Park, 1997). Most of these efforts are very recent, and much is still unknown about the relationship between implicit prejudice measures, judgmental and behavioral responses to group members, and the explicit prejudice measures provided by a variety of self-report scales.

The goal of the current paper is to propose a new measure of prejudice that involves a simple paper-and-pencil task, but which nevertheless circumvents the "willing and able" problem associated with self-report. In order to accomplish this goal, we borrowed a procedure from research on the Linguistic Intergroup Bias (LIB) by Maass and her colleagues (e.g., Maass, Milesi, Zabbini, & Stahlberg, 1995; Maass, Salvi, Arcuri, & Semin, 1989). In Maass et al.'s (1995) experiments, subjects were presented with drawings depicting members of different groups engaging in stereotype-congruent and -incongruent behaviors. Subjects were then asked to select one of four statements that best described the action in the drawing. The four descriptive statements accompanying each drawing were developed according to Semin and Fiedler's (1988) Linguistic Category Model and were intended to represent different levels of abstraction. Ranging from least to most abstract, *descriptive action verbs* refer to objective descriptions of observable behaviors that have a clear beginning and end (e.g., hit), *interpretive action verbs* describe a general class of behaviors and have positive or negative connotations (e.g., hurt), *state verbs* refer to enduring states without a clear beginning or ending (e.g., hate), and *adjectives* describe highly abstract personal dispositions (e.g., is violent).

Maass et al. found that people use relatively concrete descriptors when referring to stereotype-incongruent behaviors but relatively abstract descriptors when referring to stereotype-congruent behaviors. Thus, for example, if an Asian-American woman were seen striking someone, people would describe the event in terms of concrete behaviors ("Serena Chu hit someone"), whereas if an African-American were engaging in this same behavior, people would describe the event in abstract trait terms ("Jamal Washington is violent"). In this manner, people make no inferences from behaviors that are contrary to their expectations, but make and communicate inferences from behaviors that are congruent with their expectations.

Maass et al.'s (1989, 1995) research suggests that the LIB might play an important role in maintaining prejudice by facilitating the communication of stereotypes to others. Furthermore, the LIB has been found to mediate expectancy maintenance within a single individual (Karpinski & von Hippel, 1996) and to be exacerbated in the presence of threat (Maass, Ceccarelli, & Rudin, 1996). Thus, the LIB might be both cause and consequence of prejudice toward other groups. Finally, the LIB also seems to operate outside of awareness, as people show the LIB even when they do not show ingroup bias on more direct tasks (Franco & Maass, 1996). In conjunction, these results suggest that the LIB might be ideal as

an implicit marker of prejudice, in that those who show the LIB should tend to be high in prejudice and those who do not show the LIB should tend to be low in prejudice.

Hypotheses

The first hypothesis of these experiments is that an implicit prejudice measure based on the LIB should predict judgments of individual group members. Specifically, for the reasons stated above, people who show the LIB should evaluate group members more stereotypically than people who do not show the LIB. In contrast, because self-report measures of prejudice are plagued by the willing and able problem, such explicit measures are less likely to be useful in predicting judgments of group members.

The second hypothesis concerns the relationship between the LIB and self-report measures of prejudice. At this point, research on implicit and explicit prejudice suggests that the two types of measures tend to be dissociated (Banaji & Greenwald, 1995; Banaji & Hardin, 1996; Greenwald & Banaji, 1995). For example, Fazio et al. (1995) developed an implicit prejudice measure based on reaction times to report positive words as good or negative words as bad when the valenced words were primed by a photograph of a Caucasian or African-American. Fazio et al. found that scores on this priming measure predicted people's friendliness in an interaction with an African-American confederate, but were uncorrelated with scores on the Modern Racism Scale (MRS; McConahay, Hardee, & Batts, 1981). On the other hand, Wittenbrink et al. (1997) adapted this reaction time measure by using race-related words rather than faces as primes, and by having people make lexical decisions rather than valence decisions about the target words, and found that the measure was highly correlated with scores on explicit measures such as the MRS. Thus, even in cases in which implicit measures are methodologically and conceptually highly related, different patterns of relationships have occasionally emerged between implicit and explicit measures (see also Hense, Penner, & Nelson, 1995).

It is our position that these inconsistencies may have more to do with the nature of self-report measures than with differences among implicit measures. Self-report measures of prejudice are likely to vary in utility as a function of the environment in which they are given. When participants are concerned about social desirability, when they do not have conscious access to their prejudices, or when the environment is one in which political correctness is an important concern, self-report measures are likely to perform rather poorly. On the other hand, self-report measures might predict well in circumstances that minimize such concerns.¹ For these reasons, it becomes difficult to predict the relationship

¹ In our own laboratory, the MRS has successfully predicted memory for stereotype-congruent information (Sekaquaptewa & von Hippel, 1994) and, in contrast to Fazio et al. (1995), has also predicted the friendliness of an interaction with an African-American confederate. These results, have been difficult to replicate, however, and our informal analysis suggests that the MRS may be more effective early rather than late in the quarter, by which time subjects have become more sophisticated.

between the LIB and self-report measures because it is difficult to predict whether self-report measures will be valid or reliable. Nevertheless, given the preponderance of evidence for a dissociation between measures and the importance of the "willing and able" problem for self-report measures, we take the tentative position that the LIB will be only modestly related to self-report measures of prejudice, if at all.

In order to address these two hypotheses, we report three experiments. In the first two experiments, a self-report measure of prejudice and an implicit measure of prejudice based on the LIB were used to predict reactions to an individual African-American or Caucasian. In the third experiment, an additional implicit measure of prejudice was developed, and the relationships between this measure, the LIB measure, and a self-report measure were assessed. Although we had no particular theory concerning which self-report measures would be most appropriate in these experiments, we chose to rely on the MRS to measure prejudice toward African-Americans and the Attitude Toward Women Scale (AWS; Spence, Helmreich, & Stapp, 1973) to measure prejudice toward women. In large part, these decisions were based on the fact that these scales are among those that are most commonly used. Furthermore, the findings of Wittenbrink et al. (1997) suggest that a single common factor underlies responses on many of these measures, and thus the specific self-report measure used is unlikely to be an important factor in these experiments.

EXPERIMENT 1

Method

Subjects. Subjects were 190 Caucasian undergraduates at Ohio State University who participated in partial fulfillment of course requirements.

Materials. Evaluations of an African-American vs a Caucasian individual were assessed through the presentation of a videotaped interaction between two men. The video showed interleaved clips of one man walking down a street and another man standing on a street corner watching passersby. As the video progresses, the walking man approaches the target person standing on the street corner. As the walking man is about to pass the target, the target steps up to him and says, "Hey, I need some money. Gimme some money." The video then shows the walking man (henceforth referred to as the "requestee") look at the target (henceforth referred to as the "requester") for a moment, reach into his pocket, and give the requester an unidentifiable amount of paper money. The videotape ends as both men are seen walking in different directions. In order to manipulate the race of the requester, two different videotapes were used, but the scenes featuring the requestee were always the identical footage interleaved with footage of the requester. Half of the subjects were randomly assigned to see an African-American man ask a Caucasian man for money, and half saw a Caucasian man ask a Caucasian man for money.

Implicit prejudice was assessed with stimulus materials adapted from the Maass et al. (1989, 1995) procedure. A booklet of seven ersatz newspaper articles was assembled, of which the first, third, and fifth stories were fillers. The second story was about a basketball slam-dunk contest winner; the fourth story was about a white-collar embezzler employed by a computer firm; the sixth story was about a spelling-bee winner who was planning to attend MIT to become a physicist; and the seventh story was about a jewelry thief who lived in subsidized housing in a run-down area of the city. Stereotype congruency for each of these articles was manipulated by pairing them with photographs cut from newspapers of either an African-American or a Caucasian. The slam-dunk winner and jewelry-thief stories were stereotype-congruent when paired with an African-American face, and the spelling-bee

and embezzler stories were stereotype-incongruent when paired with an African-American face (cf. Devine & Elliot, 1995).

Race of the target was partially counterbalanced such that the article-photograph pairings were presented in two conditions. Half of the subjects were randomly assigned to read articles about an African-American slam-dunk winner and a Caucasian jewelry thief and an African-American embezzler and a Caucasian spelling-bee winner. The other half of the subjects read articles about a Caucasian slam-dunk winner and an African-American jewelry thief and a Caucasian embezzler and an African-American spelling-bee winner. In this way all subjects read one stereotype-congruent article paired with a photograph of an African-American target and one stereotype-incongruent article paired with a photograph of an African-American target. Additionally, because valence was manipulated orthogonal to stereotypicality (i.e., the slam-dunk and spelling-bee winners were described in positive terms and the jewelry thief and embezzler were described in negative terms), all subjects read both positively and negatively valenced stereotypical and counterstereotypical articles.

Below each article was a series of four statements describing the article with differing degrees of abstraction. Following Maass et al.'s (1989) use of Semin and Fiedler's (1988) linguistic category model, each statement describing the article corresponded to one of four levels of abstraction outlined above. For example, following the article about the slam-dunk winner, the four statements read: "Johnson performs 360-degree slam-dunk," "Johnson wins slam-dunk contest," "Johnson plays basketball well," and "Johnson is athletic." Thus, the sentences ranged from concrete to abstract, with the first sentence always the most concrete and the last always the most abstract. Subjects were instructed to evaluate, using 10-point scales anchored by "describes very poorly" and "describes very well," the degree to which each statement described the article presented above.

Procedure. Upon entering the laboratory, subjects were asked to watch a videotape, with the goal of recalling information from it later. After subjects watched the videotape, they were presented with a brief questionnaire ostensibly assessing their memory for the events they had seen. Included among several filler items was a question asking subjects to indicate the extent to which they thought the requester was threatening. Responses were given on an 11-point scale anchored by "not at all threatening" and "very threatening."

Subjects were then given the booklet of ersatz newspaper articles. Instructions for the booklet were intended to give the impression that the articles were culled from newspapers around the country and that the experiment was concerned with how people perceive different types of events. Subjects were instructed to read each article and indicate how well each of the four statements described the main point of the article. Upon completion of the booklet, subjects filled out the MRS, which was scored on 6-point scales anchored by "strongly agree" and "strongly disagree." The order of presentation for the explicit and implicit prejudice measures was not counterbalanced, due to the reactivity of the MRS (Dunton & Fazio, 1997; Fazio et al., 1995). Finally, subjects were debriefed, thanked, and dismissed.

Results

Classification of the prejudice measures. In order to classify subjects as high or low in implicit prejudice, two different composite scores were created using subjects' responses to the most abstract level of descriptions (i.e., adjectives) of the stereotype-congruent and -incongruent articles. We hypothesized that the propensity to process information in stereotypic ways would be associated with greater abstraction from the stereotype-congruent articles when they were paired with an African-American rather than a Caucasian and less abstraction from the stereotype-incongruent articles when they were paired with an African-American rather than a Caucasian. This focus on abstract rather than concrete descriptors was consistent with the results of Karpinski and von Hippel (1996), who found

with similar materials that the LIB only emerged with subjects' responses to the most abstract descriptions.²

For the stereotype-congruent articles (slam-dunk and jewelry thief), an implicit-prejudice score was computed by subtracting subjects' rating of the abstract description of the article paired with a Caucasian from their rating of the abstract description of the article paired with an African-American. Positive numbers on this measure indicate a greater preference for the abstract description of the article when the target was African-American rather than Caucasian. Negative numbers indicate a greater preference for the abstract description of the article when the target was Caucasian rather than African-American. A score of zero indicates no preferences for the abstract description as a function of race of the target. Because implicit prejudice was defined as the tendency to show the LIB, subjects with composite scores greater than zero were considered high in implicit prejudice, and subjects with scores less than or equal to zero were considered low in implicit prejudice.³ In this experiment, this procedure resulted in the classification of 104 subjects as low in prejudice and 85 subjects as high in prejudice.⁴

For the stereotype-incongruent articles (spelling bee and embezzler), an implicit-prejudice score was computed by subtracting subjects' rating of the abstract

² In support of this reasoning, incorporating subjects' evaluations of the less abstract descriptions into the implicit prejudice measure (by averaging evaluations of the two most abstract statements and reverse scoring and averaging evaluations of the two most concrete statements) had no effect on the pattern of results in Experiment 1, but decreased the predictive ability of the implicit prejudice measures in Experiments 2 and 3.

³ Although subjects in the current research were divided into high- and low-prejudice in a dichotomous fashion, one could imagine an analytic strategy in which the prejudice scale was treated as a continuous variable and regression analyses were used. In the current research this strategy proved to be somewhat less reliable than the dichotomization reported herein. Specifically, in Experiment 1 there were no differences in the results as a function of whether the scales were treated in a dichotomous vs a continuous fashion, in Experiment 2 the continuous scale revealed results similar to the dichotomous scale with one dependent variable and weaker results with the other, and in Experiment 3 the continuous scale revealed slightly weaker results than the dichotomous scale. Additionally, trichotomizing the prejudice scale into people with positive scores, people with scores of zero, and people with negative scores revealed in all three experiments that those with scores of zero responded to the targets in a virtually identical fashion to those with negative scores. Thus, it seems to be the case that this particular implicit prejudice scale behaves primarily in a dichotomous fashion, and that people who show no evidence of a linguistic bias (i.e., a zero score) are theoretically and empirically most similar to people who show a bias in the counter-stereotypic direction (i.e., a negative score).

⁴ The prejudice scale was also analyzed separately as a function of whether the African-American target was associated with the positive (i.e., basketball) or negative (i.e., jewelry thief) story. This analysis revealed that a somewhat larger percentage of subjects was categorized as low in prejudice when the African-American was associated with the basketball story (59%) than when he was associated with the jewelry-thief story (49%). Such a difference between the two conditions in the ratio of low- to high-prejudice subjects suggests that measurement idiosyncrasies exist in our operationalization of implicit prejudice. Nevertheless, in none of the three experiments did the counterbalancing of race (or gender) with story show main effects or interactions with the other independent variables (all F 's < 1, ns). For this reason, all analyses are reported after collapsing across the counterbalancing conditions.

description of the article paired with the African-American face from their rating of the abstract description of the article paired with the Caucasian face. As with the stereotype-congruent measure, subjects with difference scores greater than zero were considered high in implicit prejudice and subjects with difference scores less than or equal to zero were considered low in implicit prejudice. Note that whereas the stereotype-congruent measure taps subjects' increased willingness to make abstract inferences from stereotype-congruent articles, the stereotype-incongruent measure taps subjects' unwillingness to make abstract inferences from stereotype-incongruent articles. In this experiment, this procedure resulted in the classification of 121 subjects as low in prejudice and 67 subjects as high in prejudice.

Because responses to the MRS were provided on a scale that ranged from 1 to 6, the possible range of scores was from 7 to 42. Subjects were divided into high and low explicit prejudice via a median split on the MRS (Median = 16.00, range = 7 to 40, $SD = 6.2$).⁵ Consistent with predictions, the stereotype-congruent and stereotype-incongruent implicit prejudice scores were both uncorrelated with the MRS ($r = -.05$, $N = 190$, $p > .50$, $r = .07$, $N = 189$, $p > .35$, respectively), although somewhat surprisingly, they were uncorrelated with each other as well ($r = -.01$, $N = 189$, $p > .85$).

Evaluations. In order to determine whether implicit and explicit prejudice scores would predict differential responding to the African-American and Caucasian requester, subjects' ratings of threat were analyzed in two-way between-subjects ANOVA's.⁶ The first analysis examined subjects' ratings of threat as predicted by race of requester and stereotype-congruent implicit prejudice. This ANOVA revealed no main effect for implicit prejudice ($F < 1$, ns) but a marginal main effect for race of requester, such that the Caucasian requester was rated as slightly more threatening ($M = 4.18$) than the African-American requester ($M = 3.68$), $F(1, 185) = 3.30$, $p < .08$. More importantly, this main effect was qualified by a two-way interaction between race of requester and stereotype-congruent implicit prejudice, $F(1, 185) = 5.62$, $p < .02$. Simple effects analyses revealed that low implicit-prejudice subjects rated the Caucasian requester as more threatening than the African-American requester, $F(1, 185) = 8.64$, $p < .01$, whereas high implicit-prejudice subjects did not differ in their evaluations of the

⁵ Subjects were also divided into high and low explicit prejudice on the MRS (and the AWS in Experiment 3) using the ratios of high- and low-prejudice responses on the implicit prejudice measures. These changes in the classification of high- and low-prejudice on the explicit measures had no impact on their predictive ability in any of the three experiments. Additionally, treating the explicit measures as continuous variables and relying on regression techniques also failed to yield any differences in their predictive ability in Experiments 1 or 2, but did reveal a significant relationship between the AWS and the attributionally based implicit prejudice measure in Experiment 3.

⁶ This strategy of relying on separate two-way ANOVA's was adopted because three-way and four-way ANOVA's often resulted in highly unequal cell numbers. Furthermore, with the exception of a three-way and a four-way interaction that emerged with stereotype-incongruent implicit prejudice, both of which were uninterpretable, no three-way or four-way interactions emerged in any of the experiments (all F 's < 2.30 , p 's $> .10$).

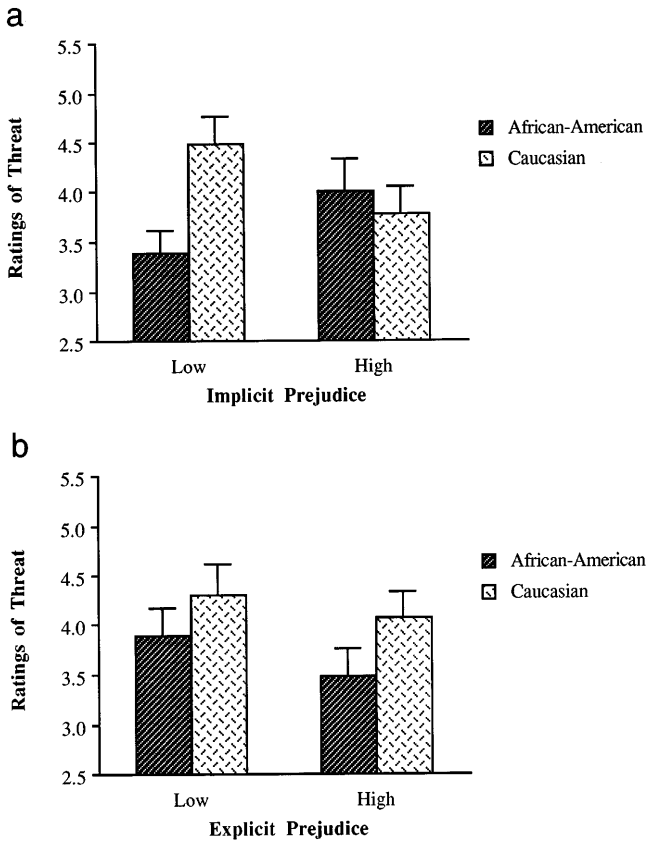


FIG. 1. (a) Subjects' ratings of threat as a function of stereotype-congruent implicit prejudice and race of the requester. (b) Subjects' ratings of threat as a function of explicit prejudice and race of the requester.

African-American and Caucasian requester, $F(1, 185) = .29, p > .55$ (see Fig. 1a).

Analyses failed to reveal main effects or interactions for explicit prejudice, $F(1, 185) = .10, p > .75$ (see Fig. 1b), or stereotype-incongruent implicit prejudice, $F(1, 184) = .07, p > .75$.

Discussion

This experiment revealed that stereotype-congruent implicit prejudice successfully predicted differential evaluations of a Caucasian and an African-American target. Subjects low in implicit prejudice evaluated the African-American requester as less threatening than the Caucasian requester, whereas subjects high in implicit prejudice showed a nonsignificant reversal of this effect. Furthermore, the explicit prejudice measure failed to predict differences in evaluations of the

Caucasian and African-American. Somewhat surprisingly, stereotype-incongruent implicit prejudice also did not predict evaluations. Although it is not clear why this measure did not work similarly to the stereotype-congruent measure, one possibility lies in the opposing nature of the task demands. Stereotype-congruent implicit-prejudice scores reflect subjects' willingness to make abstract inferences from stereotype-congruent articles, whereas stereotype-incongruent implicit-prejudice scores reflect subjects' unwillingness to make abstract inferences from stereotype-incongruent articles. As the absence of a correlation between these two measures suggests, it may simply be the case that these two tasks reflect different underlying processes. Nevertheless, it seemed worthwhile to include both types of implicit-prejudice measures in a second experiment.

Another somewhat surprising finding from Experiment 1 was that differences emerged in the evaluations of the Caucasian and African-American target among low implicit-prejudice subjects, but not among high implicit-prejudice subjects. According to theory, one would expect the opposite finding, whereby people high in prejudice would show a bias against members of a stereotyped group and people low in prejudice would show neither a bias for nor against members of a stereotyped group (although low-prejudice individuals might "bend over backwards" in their efforts to be fair and thereby overcorrect their judgments; cf. Wegener & Petty, 1995). It seems likely, however, that differences emerged only in the evaluations of low-prejudice subjects in Experiment 1 because of the direct nature of the threat question that was asked. That is, despite the fact that the threat question was embedded in a series of filler questions, it must have been relatively clear to subjects who were exposed to an African-American requester that the threat measure was assessing stereotypic responses to African-Americans. Support for this possibility can be seen in the marginal main effect that emerged in this experiment, whereby the Caucasian target was rated as more threatening than the African-American target. This finding suggests that subjects may have been reporting that the African-American was less threatening than they actually perceived him to be.

To the extent that social desirability pressures caused subjects to reduce their threat ratings of the African-American but not the Caucasian, the probability increases that significant differences will emerge in the evaluations of low-prejudice subjects, and the probability decreases that significant differences will emerge in the evaluations of high-prejudice subjects (see Fig. 2). Thus, it seems possible that differences emerged among low-prejudice but not high-prejudice subjects because of the direct nature of the threat question. Consistent with such a possibility, Wittenbrink et al. (1997) found that Caucasian subjects will even show slowed reaction times in a lexical decision-making task with stereotypic words that follow race-related words when the race-related words are presented supra-lingually, but not when the race-related words are presented subliminally. According to this interpretation of the findings of Experiment 1, shifting to an implicit dependent measure should eliminate the effect caused by social desirability and

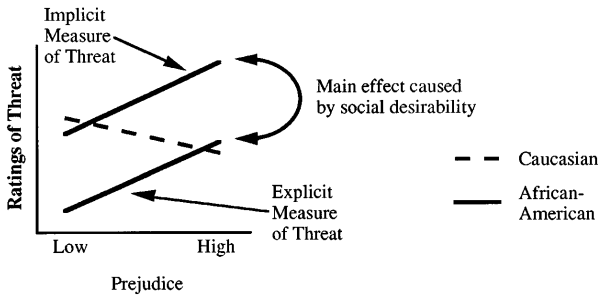


FIG. 2. The impact of social desirability on the shape of the interaction between prejudice level and race of the target.

thereby increase the probability that differences will emerge in the evaluations of high-prejudice subjects.

EXPERIMENT 2

The goal of Experiment 2 was to determine whether implicit prejudice could predict implicit evaluations of the African-American and Caucasian from the same videotape used in Experiment 1. Two different types of dependent variables were used to assess implicit evaluations: First, subjects were asked to evaluate the meekness of the requestee. We hypothesized that high implicit-prejudice subjects would rate the requestee as more meek when the requester was Caucasian than when he was African-American. In order to explain the logic behind this prediction, it is helpful to consider the following example: Suppose we see a thief holding a gun and demanding money from an unarmed passerby. We would not consider this passerby to be a very meek person if he were to give up his money. If, however, the passerby were asked for his money by a person without a weapon, we would be more likely to consider his giving up the money a rather meek act. By asking subjects about the meekness of the requestee, we can make inferences about their impressions of the requester—if it was meek to give up the money, then the requester was perceived as not threatening; if it was not meek to give up the money, the requester must have been threatening.

Second, and consistent with the cover story that we were measuring memory for the videotaped interaction, subjects were asked if they remembered whether the requester blocked the path of or touched the requestee (in actuality, the requester did neither). Blocking a person's path or touching a person while requesting money are two potentially threatening acts, but asking subjects whether they remember these behaviors should be a relatively subtle, or implicit, measure of how threatening they perceived the requester to be.

We hypothesized that, consistent with Experiment 1, implicit prejudice would predict subjects' differential responses to African-Americans and Caucasians on these implicit measures of threat. Furthermore, with these implicit-threat measures it was predicted that differences would emerge among high implicit-

prejudice subjects in their evaluations of African-Americans and Caucasians. As in Experiment 1, it was unclear whether explicit prejudice would also predict subjects' responses on these implicit-threat measures.

Method

Subjects. Subjects were 119 Caucasian undergraduates at Ohio State University who participated in partial fulfillment of course requirements.

Materials. The videotape and measures of implicit and explicit prejudice from Experiment 1 were again used in Experiment 2.

Procedure. Subjects watched the videotape and completed the implicit- and explicit-prejudice measures as in Experiment 1. The dependent measures in Experiment 2, however, were changed slightly. After a series of filler questions in which subjects were asked to recall various events from the videotape, they were instructed to evaluate the meekness of the requestee on an 11-point scale anchored by "not at all meek" and "very meek." Subjects were then asked to recall whether the requester blocked the path of the requestee and to rate how confident they were in this recollection on an 11-point scale anchored by "not at all confident" and "very confident." Additionally, subjects were asked to recall whether the requester touched the requestee and to rate how confident they were in this recollection. Subjects were then debriefed, thanked, and dismissed.

Results

Implicit-prejudice scores were computed as in Experiment 1. Additionally, subjects were again divided into high and low explicit-prejudice groups via a median split on the MRS (Median = 18.00, range = 7 to 39, $SD = 6.2$). In this experiment, the stereotype-congruent implicit-prejudice measure resulted in the classification of 69 subjects as low in prejudice and 49 subjects as high in prejudice. The stereotype-incongruent implicit-prejudice measure resulted in the classification of 72 subjects as low in prejudice and 47 subjects as high in prejudice. As in Experiment 1, no significant correlations emerged between either stereotype-congruent or stereotype-incongruent implicit prejudice and the MRS ($r = -.03$, $N = 118$, $p > .70$, $r = .09$, $N = 118$, $p > .30$, respectively), and the two implicit prejudice measures were again uncorrelated as well ($r = .15$, $N = 119$, $p > .10$).

Evaluations. Subjects' responses to the questions about whether the requester touched or blocked the path of the requestee were coded as +1 if they recalled that the requester did touch or block the requestee, and -1 if they recalled that the requester did not touch or block the requestee. Subjects' associated confidence ratings were then multiplied by these scores, and a value of 1 was added to the result if it was negative in order to create a continuous scale (see Manis, Paskewitz, & Cotler, 1986). This transformation resulted in a response scale with a positive score (from 1 to 11) indicating subjects' confidence that the requester did touch or block the path of the requestee, and 0 or a negative score (from -10

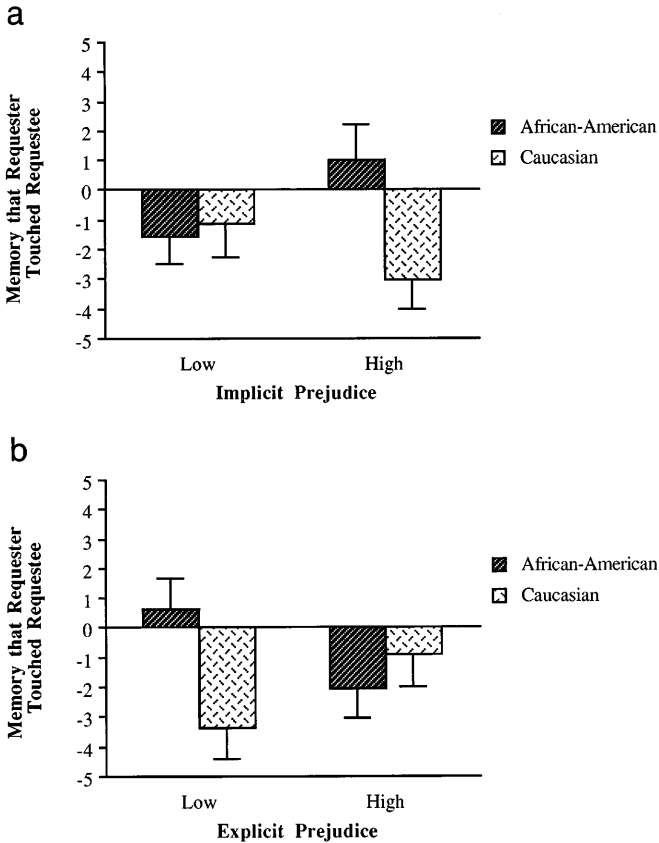


FIG. 3. (a) Subjects' recall that the requester touched the requestee as a function of stereotype-congruent implicit prejudice and race of the requester. (b) Subjects' recall that the requester touched the requestee as a function of explicit prejudice and race of the requester.

to 0) indicating subjects' confidence that the requester did not touch or block the requestee.

These "touch" and "block" variables were then analyzed in separate two-way ANOVA's. Although no main effects or interactions emerged with the block variable (F 's < 1.16 , p 's $> .25$), with the touch variable the predicted two-way interaction emerged between race of requester and stereotype-congruent implicit prejudice, $F(1, 114) = 4.29$, $p < .05$. Simple effects analyses revealed that high implicit-prejudice subjects recalled that the African-American requester was more likely than the Caucasian requester to have touched the requestee, $F(1, 114) = 5.95$, $p < .02$, whereas low implicit-prejudice subjects showed no differences in memory as a function of the race of the requester, $F(1, 114) = .11$, $p > .70$ (see Fig. 3a). No main effects emerged from this analysis (F 's < 1.87 , p 's $> .15$).

Somewhat surprisingly, the opposite two-way interaction emerged between

race of the requester and explicit prejudice, $F(1, 114) = 5.99, p < .02$. Simple effects analyses revealed that low explicit-prejudice subjects recalled that the African-American requester was more likely than the Caucasian requester to have touched the requestee, $F(1, 114) = 7.21, p < .01$ (see Fig. 3b), and no differences emerged in the memory of high explicit-prejudice subjects, $F(1, 114) = .61, p > .40$. No main effects emerged from this analysis (F 's $< 1.79, p$'s $> .15$), and no main effects or interactions emerged with stereotype-incongruent implicit prejudice (F 's $< 1.75, p$'s $> .15$).

Turning to subjects' meekness ratings, a two-way interaction emerged between race of requester and stereotype-congruent implicit prejudice, $F(1, 113) = 5.54, p < .03$. Simple effects analyses revealed that low implicit-prejudice subjects rated the requestee as more meek when the requester was African-American than when he was Caucasian, $F(1, 113) = 4.01, p < .05$. Thus, subjects low in implicit prejudice apparently perceived the Caucasian requester as more threatening than the African-American requester. High implicit-prejudice subjects, however, showed a nonsignificant reversal of this effect: the requestee was evaluated as slightly more meek when the requester was Caucasian than when he was African-American, $F(1, 113) = 1.96, p = .16$. Although not statistically significant, this difference suggests that high implicit-prejudice subjects may have perceived the African-American requester as slightly more threatening than the Caucasian requester. Thus, stereotype-congruent implicit prejudice predicted subjects' meekness evaluations of the requestee in a manner largely consistent with predictions (see Fig. 4a).

As with the touch variable, the opposite two-way interaction emerged between race of requester and explicit prejudice, $F(1, 113) = 5.41, p < .03$. Simple effects analyses revealed that low-prejudice subjects did not differ in their evaluations of the meekness of the requestee as a function of the race of the requester, $F(1, 113) = 1.40, p > .20$. High-explicit-prejudice subjects, however, evaluated the requestee as more meek when he was approached by an African-American requester than when he was approached by a Caucasian requester, $F(1, 113) = 4.39, p < .04$ (see Fig. 4b). This finding suggests that high-explicit-prejudice subjects perceived the Caucasian as more threatening than the African-American. No other main effects or interactions emerged from this analysis, and no effects emerged for stereotype-incongruent implicit prejudice (all F 's $< 1, ns$).

Finally, in order to assess the extent to which the two implicit evaluation measures were tapping the same construct, subjects' meekness ratings and the touch variable were correlated with one another. As predicted, subjects' ratings on the two implicit-threat variables were negatively correlated ($r = -.25, N = 116, p < .01$).

Discussion

Experiment 2 provides additional support for the notion that implicit prejudice predicts differential evaluations of individual African-Americans and Caucasians. Subjects' ratings on two implicit-threat measures were predicted by their stereotype-

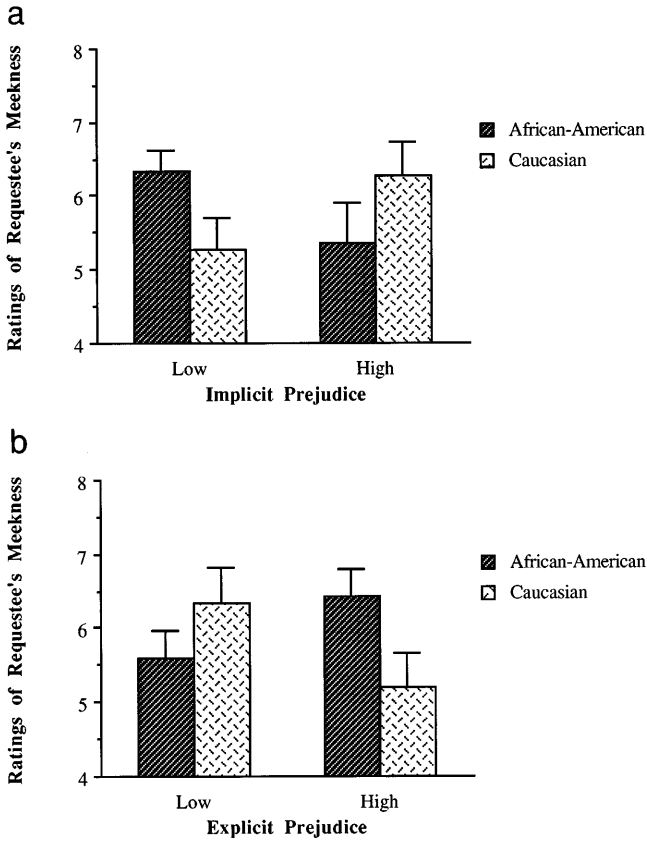


FIG. 4. (a) Subjects' ratings of the requestee's meekness as a function of stereotype-congruent implicit prejudice and race of the requester. (b) Subjects' ratings of the requestee's meekness as a function of explicit prejudice and race of the requester.

congruent implicit-prejudice scores, indicating that high-prejudice subjects perceived the African-American as more threatening than the Caucasian and low-prejudice subjects showed a reversal of this effect. Furthermore, the correlation between the two implicit dependent measures suggests that they are tapping the same construct, despite the rather large discrepancy in the type of information that was asked by the two implicit-threat variables. Finally, as in Experiment 1, explicit prejudice and stereotype-incongruent prejudice did not yield the predicted differences in evaluations of individual Caucasians versus African-Americans.

Although it may seem that Experiments 1 and 2 are inconsistent with each other with regard to whether differences emerged in the evaluations of those who are low or high in implicit prejudice, as we proposed earlier, this inconsistency may be more illusory than substantive. Congruent with the notion that social desirability pressures have an impact on explicit but not implicit measures, subjects rated

the Caucasian requester as more threatening than the African-American requester on the explicit measure in Experiment 1, but not on the implicit measures in Experiment 2. It was this main effect that caused differences to emerge only among low-prejudice subjects in Experiment 1, and its absence that caused differences to emerge among high-prejudice subjects in Experiment 2 (see Fig. 2).

The results of Experiments 1 and 2 provide evidence that implicit prejudice toward African-Americans is distinct from explicit prejudice and is an effective predictor of evaluations of individual African-Americans. The goal of Experiment 3 was to extend these findings by developing an additional measure of implicit prejudice, based on the same principle of assessing information-processing biases, and examine whether such a measure would be related to the implicit-prejudice measure based on the LIB. In order to test this possibility, Experiment 3 examined prejudice toward women.

EXPERIMENT 3

In Experiment 3 one implicit prejudice measure was again based on the LIB, with a new set of articles that were written to be consistent with the female stereotype. Because Experiments 1 and 2 failed to show any effects for stereotype-incongruent implicit prejudice, Experiment 3 only contained a measure of stereotype-congruent implicit prejudice. A second implicit-prejudice measure relied on the assessment of biased attributional responding. In particular, in Experiment 3 we examined the likelihood that subjects would engage in attributional reasoning when they encountered stereotype-congruent and incongruent behaviors from males and females. As Hastie (1984) has demonstrated, people are more likely to explain behaviors when they are incongruent with expectancies than when they are congruent with expectancies. Because this explanatory processing renders the seemingly incongruent behavior no longer incongruent, explanatory processing should serve the same function as the LIB in maintaining stereotypes in the face of incongruency. Thus, this attributional bias may also be an effective measure of implicit prejudice. The goal of Experiment 3 was to determine whether a measure of implicit prejudice based on the LIB (henceforth referred to as linguistic implicit prejudice) would be correlated with a measure of implicit prejudice based on the tendency to engage in explanatory processing when confronted with stereotype-incongruency (henceforth referred to as attributional implicit prejudice).

Method

Subjects. Subjects were 156 male undergraduates at Ohio State University who participated in partial fulfillment of course requirements.

Materials. Linguistic implicit prejudice was assessed with ersatz newspaper articles in a procedure similar to that used in Experiments 1 and 2. In this experiment, however, the targets of the newspaper articles were either male or female, and this manipulation was conveyed through the use of gender-specific names such as "James" and "Molly." Two experimental articles were embedded

in a set of four newspaper articles in which the first and third articles were fillers. The second and fourth articles were counterbalanced, as was the gender of the protagonist in the articles, although both articles described a person who engaged in female-stereotypic behaviors. One of these articles described a person who won a community award for running a day-care facility and who loved to take care of children. The other article described a person who was extremely dependent on his/her spouse and who had not been able to get anything done around the house since the spouse had abandoned him/her. Thus, all subjects read one female stereotype-congruent article paired with a male name and one female stereotype-congruent article paired with a female name. Subjects then rated statements describing the articles at varying levels of abstraction, as in Experiments 1 and 2.

Attributional implicit prejudice was assessed with a set of 22 sentence beginnings. Of these sentence beginnings, 10 described gender-neutral behaviors (e.g., Alina brushed her teeth), 6 described stereotypically female behaviors (e.g., Wendy babysat the neighbors' kids), and 6 described stereotypically male behaviors (e.g., Steve studied the engineering manual). Half of the behaviors were paired with a male name and half were paired with a female name, and the specific pairings of a behavior with a name were counterbalanced between subjects. Subjects were instructed to read the sentence beginnings and continue them in any fashion that created a grammatically correct sentence (see Hastie, 1984).

Procedure. Upon entering the laboratory subjects were told that they would engage in several different tasks that were broadly concerned with impression formation. Subjects were then presented with the newspaper articles that were used to assess linguistic implicit prejudice. When they had finished rating these articles, subjects were presented with the sentence beginnings that were used to assess attributional implicit prejudice. When subjects had completed the sentence-continuation task, explicit prejudice toward women was assessed with the AWS. Upon completion of this scale, subjects were debriefed, thanked, and dismissed.

Results

Because this experiment presented subjects only with female-stereotypic articles, for the linguistic implicit-prejudice measure only a stereotype-congruent score was computed. As in Experiments 1 and 2, the linguistic implicit-prejudice score was computed by subtracting subjects' rating of the most abstract description of the article describing a male target from their rating of the most abstract description of the article describing a female target. Subjects with scores greater than zero were considered high in linguistic implicit prejudice, and subjects with scores less than or equal to zero were considered low in linguistic implicit prejudice. In this experiment, dichotomizing by this subtraction procedure resulted in the classification of 87 subjects as low in prejudice and 66 subjects as high in prejudice.

In order to compute an attributional implicit-prejudice score, subjects' sentence completions were scored by two raters as to whether the completions were explanatory or nonexplanatory in nature (raters were blind to hypotheses and

differences between raters were resolved by discussion). A difference score was created by subtracting the number of explanations that subjects provided for stereotype-congruent behaviors from the number of explanations they provided for stereotype-incongruent behaviors. Thus, the attributional prejudice measure was scored by subtracting the number of explanations for male behaviors when exhibited by a male from the number of explanations for male behaviors when exhibited by a female. This difference score was then averaged with a difference score created by subtracting the number of explanations for female behaviors when exhibited by a female from the number of explanations for female behaviors when exhibited by a male. Positive numbers on this measure indicate a greater tendency to explain behaviors when they are stereotype-incongruent than when they are stereotype-congruent.

As with the linguistic-prejudice score, this attributional score was dichotomized such that subjects who showed an average difference score greater than zero were considered high in attributional prejudice, and subjects who showed a difference score less than or equal to zero were considered low in attributional prejudice. This procedure resulted in the classification of 104 subjects as low in prejudice and 50 subjects as high in prejudice. Subjects were then divided into high and low explicit prejudice groups via a median split on the AWS (Median = 33.00, range = 19 to 53, $SD = 6.2$). Consistent with predictions, the linguistic implicit-prejudice measure was correlated with the attributional implicit-prejudice measure ($r = .21, N = 153, p < .01$), but neither the linguistic nor the attributional prejudice measure was correlated with the explicit-prejudice measure ($r = .01, N = 153, p > .90$, and, $r = .09, N = 152, p > .25$, respectively).⁷

Discussion

The results of Experiment 3 provide support for the hypothesis that different measures of implicit prejudice based on biased information-processing would be related to one another. In this experiment, a linguistically based implicit-prejudice measure toward women was correlated with an attributionally based implicit-prejudice measure. Subjects who preferred abstract descriptors when they encountered stereotype-congruent behavior also generated explanations when they encountered stereotype-incongruent behavior. Thus, this experiment extends Experiments 1 and 2 by demonstrating the convergent validity of the implicit-prejudice measure, in that an implicit-prejudice measure derived from the linguistic research of Maass et al. (1989, 1995) was related to an implicit-prejudice measure derived from the attributional research of Hastie (1984).

⁷ Recall, however, that as noted in footnote 3, when the three prejudice measures were treated as continuous scales, the AWS was significantly correlated with the attributional prejudice scale ($r = .17, N = 152, p < .04$) but not the linguistic prejudice scale ($r = .01, N = 153, p > .90$). The two implicit prejudice measures were still correlated with one another when treated as continuous scales, although as noted in footnote 3, the relationship was somewhat weaker ($r = .15, N = 153, p < .06$). This was the only case in which the correlations between implicit and explicit measures differed in significance as a function of whether they were treated as continuous or dichotomous scales.

Furthermore, these implicit-prejudice measures were not systematically related to an explicit-prejudice measure.

GENERAL DISCUSSION

The results of these experiments provide support for the hypothesis that implicit measures of prejudice based on biased processing should predict reactions to individual group members. In Experiment 1, implicit prejudice toward African-Americans predicted subjects' differential responses to African-American and Caucasian targets. Experiment 2 replicated this finding with indirect evaluations of African-American and Caucasian targets. Experiment 3 extended these findings by demonstrating that an implicit measure of gender prejudice based on biased linguistic processing was correlated with an implicit measure of gender prejudice based on biased attributional processing.

The results of these experiments also provide support for the hypothesis that implicit and explicit measures of prejudice should be largely dissociated. In all three experiments implicit and explicit measures of prejudice were uncorrelated. In this regard the current data are closer to the findings of Banaji and her colleagues (e.g., Banaji & Greenwald, 1995; Banaji & Hardin, 1996) and Fazio et al. (1995) than they are to those of Wittenbrink et al. (1997) and Hense et al. (1995). Whether these inconsistencies across laboratories are a function of the type of implicit measures, the type of environment in which the explicit measures are given, or some interaction between the two remains to be seen.

Finally, it should be noted that although it is likely that the implicit measures were more successful than the explicit measures because they avoid the "willing and able" problem, it is also possible that other differences between the measures were critical in producing the results that emerged. One potentially important difference is that the explicit measures were designed to tap the *content* of people's prejudicial attitudes, whereas the implicit measures were designed to tap the prejudicial *processes* in which they engage (von Hippel, Sekaquaptewa, & Vargas, 1995). It may well be the case that prejudice contains important processing components and that the predictive ability of prejudice measures is enhanced to the degree that they tap cognitive processes.

One might go one step further in this analysis and propose that measures of prejudice that assess cognitive processes might even be considered *procedural* measures of prejudice, as they may tap habitual responses to members of stereotyped groups. Such a possibility is consistent with many prior conceptions of prejudice, in which notions of habitual responses and negative attitudes are often intermixed. For example, Brewer (1994) defines prejudice as, "a personal disposition *or response orientation* toward a particular social group" (p. 317, emphasis added). Because a response orientation is a proclivity to respond in a particular way, this definition includes the notion that prejudice is partially composed of biases in information processing. In a related vein, Devine (1989) has argued that becoming nonprejudiced "can be likened to the breaking of a bad *habit*" (p. 15, emphasis added). This notion of prejudice as habit is tacit

recognition that prejudice has procedural components, as procedural knowledge is defined as skills and habits.

Nevertheless, whether the implicit measures reported herein are effective because they tap processing biases or because they operate outside of awareness, the utility of these measures suggests that a variety of similar measures might be worth developing that rely on other cognitive biases that have been documented in the literature. For example, people have been shown to maintain their stereotypes through biased attributional processing (Pettigrew, 1979), biased memorial processing (Fyock & Stangor, 1994), and biased assimilative processing (Hilton & von Hippel, 1990). Implicit-prejudice measures could be adapted from the tendency to show any or all of these biases. The underlying idea behind such a conception of implicit prejudice is that it will manifest itself in a variety of specific biases that are associated with one another by virtue of their relationship to the general tendency to process information in stereotype-maintaining ways.

Conclusions

The results of these experiments provide further evidence for the utility of measuring both implicit and explicit components of prejudice. The measures used to assess implicit prejudice are distinctly different from commonly used explicit measures of prejudice, and yet were effective at predicting differential responses to individual members of stereotyped groups. Furthermore, these implicit measures are easily used and readily adapted from previous research. In combination with the burgeoning literature on implicit prejudice, the findings of the current research suggest that a complete understanding of prejudice can no longer be obtained simply by asking people what they think and feel about various groups.

REFERENCES

- Banaji, M. R., & Greenwald, A. G. (1995). Implicit gender stereotyping in judgments of fame. *Journal of Personality and Social Psychology*, **68**, 181–198.
- Banaji, M. R., & Hardin, C. D. (1996). Automatic stereotyping. *Psychological Science*, **7**, 136–141.
- Brewer, M. B. (1994). The social psychology of prejudice: Getting it all together. In M. P. Zanna & J. M. Olson (Eds.), *Psychology of prejudice: The Ontario symposium* (Vol. 7, pp. 315–329). Hillsdale, NJ: Erlbaum.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, **56**, 5–18.
- Devine, P. G., & Elliot, A. J. (1995). Are racial stereotypes really fading? The Princeton trilogy revisited. *Personality and Social Psychology Bulletin*, **21**, 1139–1150.
- Dovidio, J. F., Brigham, J. C., Johnson, B. T., & Gaertner, S. L. (1996). Stereotyping, prejudice, and discrimination: Another look. In C. N. Macrae, C. Stangor, & M. Hewstone (Eds.), *Stereotypes and stereotyping*. pp. 276–319. New York: Guilford.
- 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., 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.
- Franco, F. M., & Maass, A. (1996). *The linguistic intergroup bias: Is it under intentional control?* Paper presented at the European Association of Experimental Social Psychology, Gmunden, Austria.

- Fyock, J., & Stangor, C. (1994). The role of memory biases in stereotype maintenance. *British Journal of Social Psychology*, **33**, 331–343.
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, **102**, 4–27.
- Hastie, R. (1984). Causes and effects of causal attribution. *Journal of Personality and Social Psychology*, **46**, 44–56.
- Hense, R. L., Penner, L. A., & Nelson, D. L. (1995). Implicit memory for age stereotypes. *Social Cognition*, **13**, 399–415.
- Hilton, J. L., & von Hippel, W. (1990). The role of consistency in the judgment of stereotype-relevant behaviors. *Personality and Social Psychology Bulletin*, **16**, 430–448.
- Hilton, J. L., & von Hippel, W. (1996). Stereotypes. In J. T. Spence, J. M. Darley, & D. J. Foss (Eds.), *Annual Review of Psychology* (Vol. 47, pp. 237–271). Palo Alto: Annual Reviews, Inc.
- Karpinski, A. T., & von Hippel, W. (1996). The role of the Linguistic Intergroup Bias in expectancy maintenance. *Social Cognition*, **14**, 141–163.
- Maass, A., Ceccarelli, R., & Rudin, S. (1996). Linguistic Intergroup Bias: Evidence for in-group-protective motivation. *Journal of Personality and Social Psychology*, **71**, 512–526.
- Maass, A., Milesi, A., Zabbini, S., & Stahlberg, D. (1995). Linguistic intergroup bias: Differential expectancies or in-group protection? *Journal of Personality and Social Psychology*, **68**, 116–126.
- Maass, A., Salvi, D., Arcuri, L., & Semin, G. (1989). Language use in intergroup contexts: The linguistic intergroup bias. *Journal of Personality and Social Psychology*, **57**, 981–993.
- Manis, M., Paskewitz, J., & Cotler, S. (1986). Stereotypes and social judgments. *Journal of Personality and Social Psychology*, **50**, 461–473.
- McConahay, J. B., Hardee, B. B., & Batts, V. (1981). Has racism declined? It depends upon who's asking and what is asked. *Journal of Conflict Resolution*, **25**, 563–579.
- Pettigrew, T. F. (1979). The ultimate attribution error: Extending Allport's cognitive analysis of prejudice. *Personality and Social Psychology Bulletin*, **5**, 461–476.
- Sekaquaptewa, D., & von Hippel, W. (1994). *Recall for stereotype-congruent and -incongruent information depends on prejudice level*. Paper presented at the Midwestern Psychological Association, Chicago, IL.
- Semin, G. R., & Fiedler, K. (1988). The cognitive functions of linguistic categories in describing persons: Social cognition and language. *Journal of Personality and Social Psychology*, **54**, 558–568.
- Sniderman, P. M., & Tetlock, P. E. (1986). Symbolic racism: Problems of political motive attribution. *Journal of Social Issues*, **42**, 129–150.
- Spence, J., Helmreich, R., & Stapp, J. (1973). A short version of the Attitudes Toward Women Scale. *Bulletin of the Psychonomic Society*, **2**, 219–220.
- von Hippel, W., Sekaquaptewa, D., & Vargas, P. (1995). On the role of encoding processes in stereotype maintenance. In M. P. Zanna (Ed.), *Advances in Experimental Social Psychology*, (Vol. 27, pp. 177–254). San Diego, CA: Academic Press.
- Wegener, D. T., & Petty, R. E. (1995). Flexible correction processes in social judgment: The role of naive theories in corrections for perceived bias. *Journal of Personality and Social Psychology*, **68**, 36–51.
- Wittenbrink, G., Judd, C. M., & Park, B. Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology*, **72**, 262–274.