Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Experimental Social Psychology 39 (2003) 177–183
Journal of Experimental Social Psychology
www.elsevier.com/locate/jesp

Effects of situational power on automatic racial prejudice

Jennifer A. Richeson[a,*] and Nalini Ambady[b]

[a] *Department of Psychological and Brain Sciences, Dartmouth College, 6207 Moore Hall, Hanover, NH 03755, USA*
[b] *Harvard University, USA*

Received 22 December 2001; revised 1 March 2002

Abstract

This study examined the influence of situational power on automatic racial prejudice. White females anticipated participating in either an interracial or same-race interaction in one of two roles: *superior* or *subordinate*. Their racial attitudes were measured via the Implicit Association Test (Greenwald, McGhee, & Schwartz, 1998). Results revealed that both the racial composition of the anticipated dyad and participants' situational roles influenced automatic racial attitudes. Specifically, whites assigned to the high-power role of a superior of a black individual revealed more racial bias than whites assigned to the lower-power role of a subordinate. By contrast, situational power had no influence on the automatic bias of whites anticipating same-race interactions. These results reveal the manner in which situational power hierarchies serve to reinforce existing social stratification. Implications for diversity efforts and attitude change are discussed.

Effects of situational power on automatic racial prejudice

Individuals in positions of power relative to others often engage in thoughts and behaviors that serve to maintain their positions of authority. For instance, powerful individuals tend to disregard individuating information about their interaction partners, compared to individuals who do not hold positions of power, relying instead on category-based information such as stereotypes (Erber & Fiske, 1984; Goodwin, Gubin, Fiske, & Yzerbyt, 2000). Because situational power is often confounded with sociocultural group status, the cognitive biases of individuals who hold positions of relative situational power also serve to maintain existing social stratification (Fiske, 1993; Jost & Banaji, 1994; Operario, Goodwin, & Fiske, 1998). That is, during interactions between members of stigmatized and non-stigmatized groups (i.e., intergroup dyads), members of dominant groups are more often in positions of power vis-à-vis members of stigmatized groups (Sidanius & Pratto, 1993). Consequently, the cognitive biases of powerful individuals serve to maintain the dominance of powerful sociocultural groups. The purpose of the present study was to examine whether the influence of situational power on cognitive biases, such as stereotyping, extends to attitudinal biases. Specifically, we examined the influence of holding situational roles of differential power for an upcoming interracial interaction on white Americans' automatic racial prejudice against black Americans. Just as power-discrepant cognitive biases support social stratification, individuals in positions of relative situational power may also reveal attitudes that are more biased than individuals in lower-power positions.

Individuals' attitudes and beliefs concerning racial and gender groups can seemingly be activated without conscious awareness of the activation (Bargh, Chaiken, Govender, & Pratto, 1992; Bargh, Chaiken, Raymond, & Hymes, 1996; Fazio, Sanbonmatsu, Powell, & Kardes, 1986; Purdue & Gurtman, 1990). Fazio, Jackson, Dunton, and Williams (1995) demonstrated, for instance, the automatic activation of racial attitudes. Specifically, white participants responded faster to negative target adjectives when they were preceded by primes that were photographs of black Americans compared to when they were preceded by photographs of white Americans. Presumably, because participants held negative attitudes towards blacks, it was relatively

* Corresponding author.
E-mail address: jriches@dartmouth.edu (J.A. Richeson).

0022-1031/02/$ - see front matter © 2002 Elsevier Science (USA). All rights reserved.
doi:10.1016/S0022-1031(02)00521-8

easy for them to process, and therefore to respond to, adjectives that were also negative (i.e., congruent with the valence of the racial prime). Differential automatic evaluations of different racial groups have also been detected using a method developed by Greenwald and his colleagues (i.e., the Implicit Association Test: IAT) (Dasgupta, McGhee, Greenwald, & Banaji, 2000; Dasgupta & Greenwald, 2001; Greenwald et al., 1998; Ottaway, Hayden, & Oakes, 2001). Specifically, white participants have been found to associate "pleasant" words and white names more readily than they associate "pleasant" words and black names.[1] The differential ease with which pleasant is associated with white names rather than black names reflects an automatic preference for the racial category "white" relative to "black"—that is, a biased racial attitude.

Despite the apparent robustness of automatic attitude biases, some recent research finds that they may be quite malleable to contextual influences. For instance, imagining an agentic woman reduced automatic gender stereotyping (Blair, Ma, & Lenton, 2001), and exposure to a black experimenter reduced whites' automatic racial bias (Lowery, Hardin, & Sinclair, 2001). Similarly, Rudman, Ashmore, and Gary (2001) found that white students taking a diversity seminar led by and black professor showed a significant decrease in racial bias at the end of the course, whereas students taking a methods course with a white professor did not. Further, Dasgupta and Greenwald (2001) found that white and Asian American participants revealed less implicit prejudice regarding blacks during an IAT if they had recently been exposed to admired blacks (e.g., Michael Jordan) and disliked whites (e.g., Jeffrey Dahmer), compared to recent exposure to disliked blacks (e.g., Mike Tyson) and admired whites (e.g., John F. Kennedy). Thus, participants' automatic racial prejudice was attenuated when assessed in a context in which they were reminded of "positive" black individuals. Quite remarkably, this effect lasted for at least 24 h, suggesting that situational influences on implicit attitudes may be relatively enduring. Taken together, this research suggests that exposure to relatively atypical exemplars of stigmatized groups (e.g., agentic women, atypical blacks) can reduce automatic prejudice and stereotyping toward those groups.

In the present study, we explored the extent to which exposure to stigmatized individuals holding a high-power situational role may also reduce automatic bias. If atypical exemplars of groups generate different attitudes than more prototypical group exemplars, then recent, accessible experiences with a member of a stigmatized social group holding an atypical situational role may also influence participants' attitudes regarding that group. Direct evidence in support of this hypothesis stems from the work examining the influence of situational power on stereotyping. As mentioned previously, powerful individuals tend to disregard individuating information about their interaction partners, relying more on category-based sources of information, compared to individuals who do not hold positions of power (Operario et al., 1998). For instance, powerful individuals pay more attention to stereotype-consistent, rather than stereotype-inconsistent, evidence regarding their interaction partners (Fiske, 1993; Goodwin et al., 2000). By contrast, individuals in lower-power positions are known to individuate their partners, presumably because they are dependent on them for important outcomes and, therefore, are motivated to pay particular attention to them (Fiske, 1993).

If the influence of situational power on automatic attitudes mirrors its impact on stereotype activation, then individuals in positions of power should reveal a greater degree of automatic prejudice than individuals in lower-power positions. Specifically, individuals in subordinate positions with black superiors should attend to available individuating information regarding their partners, rather than to category-based knowledge. Attending to individuating information regarding a black interaction partner (so long as that information is not negative) should lead participants in subordinate roles to view their black interaction partners relatively positively, in comparison to the perception of blacks in general. Furthermore, recent exposure to a positive black exemplar should lead individuals to find it easier to associate "black with pleasant" and harder to associate "black with unpleasant," during an assessment of automatic racial attitudes such as the IAT. Consequently, they are likely to reveal attitudes that are less biased. By contrast, individuals in higher-power positions with black interaction partners should ignore individuating characteristics about their partners, attending only to category-based information. Thus, their racial attitudes should be more biased than those generated by participants in lower-power roles with black interaction partners. In sum, the prediction of the present study is that white participants' racial bias in anticipation of an interracial dyadic interaction will depend on their relative situational power.

Present study

The purpose of the present study was to examine the influence of situational power-discrepant roles on implicit racial attitudes. Specifically, we were interested in whether whites' attitudes regarding blacks would differ when they held a relatively powerful role for an anticipated interaction with a black individual, compared to

[1] The stimuli used were names that had been pre-tested to be stereotypically associated with white Americans or black Americans. Photographs of whites and blacks have also been used, resulting in similar effects.

when they held a relatively low-power role for the interaction. In order to investigate this question, we assessed the automatic racial attitudes held by white females who anticipated participating in an interracial or same-race interaction. Further, participants were assigned to one of two roles—task *superior* or *subordinate*—for the upcoming interaction, and informed that their interaction partner held the complementary role. For instance, participants assigned to the superior role were led to believe that their partner had been instructed to be their subordinate. Because individuals in lower-power roles tend to individuate their interaction partners more than individuals in higher-power roles, participants in the *subordinate* role with a black interaction partner should reveal less bias, compared to participants in the *superior* role for an interracial interaction. Roles, however, were not expected to impact the attitudes of participants anticipating an interaction with another white individual.

Method

Participants and design

Forty-four white female students at a private New England College completed this experiment for partial course credit or for a monetary reward of $5. The design was a 2 (target race: black, white) × 2 (interaction role: superior, subordinate) Factorial.

Procedure and stimuli

Upon her arrival to the laboratory, each participant was greeted by a white female experimenter, escorted into a room in which there was a desk and a computer, and seated in a chair away from the computer. Then, the participant was told "this is a study on computer-based task performance in work environments. We are interested in how work environments are responding to the increased use of computers in group assignments, projects, and tasks. Thus, in the course of the study you will be working interactively on a computer task with a fellow student." Participants were led to believe that there was another student who was simultaneously being prepared to participate in the study as their task partner. After, they read and signed a consent form.

Introduction of relative power

Each participant was assigned a situational role for the task: *superior* or *subordinate*. Further, participants were led to believe that their partners were assigned the role in complement to their own. That is, participants assigned the superior role were led to believe that their partners had been assigned the subordinate role and vise-versa. Participants' relative power was directly linked to their roles by manipulating their expectations regarding evaluation (Goodwin et al., 2000; Raven, 1993). Specifically, participants in the superior role were told, "You are the superior and your partner is your subordinate. Therefore, you will be evaluating your subordinate's task performance." Conversely, participants in the subordinate role were told, "You are the subordinate and your partner is your superior. Included in the role of a subordinate is being evaluated. Therefore, your superior will be evaluating your task performance."

Introduction to partner

After roles were induced, participants were told that "it is sometimes helpful to have some information about someone before having to work with them." They were then given a "Profile Sheet" on which to indicate specific biographical characteristics, including their name, race, sex, class (in school), activities, hobbies, and interests. After the participant completed the profile sheet, a Polaroid picture was taken of her.

The experimenter then gave the participant a Profile Sheet that presumably had been completed by her interaction partner, as well as an accompanying Polaroid. The experimenter instructed the participant to review her partner's profile, and indicated that she was taking her profile and photograph to a nearby room for her partner to review. Based on their target race condition, participants were shown a Polaroid of either a black or a white female, college-age student who attended a different university (the photograph was randomly selected from photographs of two black females and two white females).[2] In addition, the Partner Profile Sheet indicated the partner's name, race, and sex, as well as activities, professional interests, and hobbies, all of which were fabricated and experimentally controlled. The inclusion of activities, professional interests, and hobbies allowed for the introduction of individuating information about the interaction partner. Specifically, the interaction partner was identified as being involved in a mentoring program, as enjoying playing the piano, and as interested in a career in law. Overall, the individuating information was expected to leave participants with a relatively positive impression of their interaction partner, as well as to indicate that she was quite typical of students at the college, and somewhat atypical of the stereotype of black Americans. Hence, the Profiles and photographs served several purposes: (1) to convince the participant that she would be interacting with another student during the study; (2) to ensure that the participant knew the race and gender of her task partner, and

[2] All participants expected to interact with a female (i.e., same-sex) student, in order to avoid the possibility that gender would introduce a dimension of sociocultural power difference, and potentially dilute the race effect.

knew that her task partner had the same information about her; (3) to reinforce the participant's role—the appropriate role was prominently marked at the top of each Profile Sheet; and finally (4) to provide participants with individuating information about their interaction partner. After a couple of minutes, the experimenter returned and asked the participant if she were ready to begin working on the computer task.

Implicit Association Test

When the participant indicated that she was ready to begin the computer task, she was told, "As you know, we are studying computer-based task performance. Thus, you and your superior (or subordinate) will be working on a computer-based task. In order for us to obtain a more accurate measure of team performance, you and your superior (or subordinate) will be given a chance to work on the task separately to get a baseline, as well as to give you both practice. Afterwards, you will work on the task together and interactively." Next, they were directed to the desk in front of a Compaq Presario microcomputer with a 17" monitor. The experimenter told the participant, "The task that you will be working on is a word categorization task. The instructions will be presented by the computer." The experimenter then began the IAT program and left the room.

The IAT is a measure of learning that assesses the ease with which categories are associated (Greenwald et al., 1998). The task involved five key phases for which all instructions and stimuli were presented by the computer. Across all phases, stimulus words were presented and participants were required to categorize the words as quickly as possible as belonging to one of four categories (i.e., White, Black, Pleasant, or Unpleasant) to which they were introduced at the start of the study. For instance, the name "Amber" would be classified as belonging to the category "White" by pressing an appropriate key (either on the left or the right side of the keyboard). The computer did not allow incorrect responses. Phases 1 and 2 were category learning phases in which participants categorized words from the Pleasant/Unpleasant sets (e.g., joy, poison) and White/Black name sets (e.g., Amanda, Tameka), respectively. In Phase 3, the first dual-categorization phase, words from all four of the categories were presented and participants had to categorize them appropriately. In Phase 4 the appropriate response keys of one of the category sets (e.g., Pleasant/Unpleasant) were reversed, and participants were presented with several trials to practice the new categorization scheme. Finally, in Phase 5, the second dual-categorization phase, again, words from all four of the categories were presented and participants had to categorize them appropriately. Thus, Phases 3 and 5 were the critical dual-categorization response phases, assessing the degree of association between categories. Each consisted of 40 trials. In one version of the task,[3] Phase 3 instructed participants to categorize the stimulus words in a manner such that white and black names were procedurally linked to the evaluative categories "Pleasant" and "Unpleasant;" respectively. In phase 5, the instructions reversed the procedural links such that the category "White" was now linked to "Unpleasant" and "Black" was linked to "Pleasant". All response latencies were recorded and saved on the computer. After the IAT, participants were debriefed, thanked, and compensated.

IAT bias

The difference between response latencies during Phase 3 and response latencies during Phase 5 has been shown to be an index of the degree to which an individual favors one racial category over the other (i.e., white American versus black American) (Greenwald et al., 1998). For instance, when the instructions require individuals to respond that a stimulus word (e.g., Amber) belongs to the category "White" or that a stimulus word (e.g., sunshine) belongs to the category "Pleasant" by pressing the same key on the keyboard (*White*+/*Black*− Phase), they complete the task faster than when they must press the same key to indicate that a stimulus word is a member of the category "White" or a member of the category "Unpleasant" (*White*−/*Black*+ Phase). If responses to stimulus words during the "White+/Black−" Phase are faster than responses during the "White−/Black+" Phase, then the responder is thought to have a biased implicit racial attitude in favor of white Americans relative to black Americans. This difference in response times during the two phases is the primary dependent variable examined in the present work.

Results

The response latencies from the critical dual-categorization phases (3 and 5) were used to assess automatic racial attitudes. Consistent with the data-trimming guidelines reported in Greenwald et al. (1998), all latencies under 300-ms and over 3000-ms were re-coded as 300- and 3000-ms, respectively. Additionally, all response latencies were log-transformed to conform to normality assumptions. Next, the log-response times associated with trials during the "White+/Black−" category pairing phase were averaged, and the log-response times associated with trials during the "White−/Black+" category pairing phase were averaged for each participant. Finally, each participant's mean response log-latency for the White+/Black− phase was subtracted from their mean response log-latency for the White−/Black+ phase, in order to index their degree of *automatic racial*

[3] The order of category pairing (White+/Black−; White−/Black+) was counter-balanced across participants.

Table 1
Mean response times for White−/Black+ Phase (W−/B+) and White+/Black− Phase (W+/B−)

Participant role	White partner			Black partner		
	W−/B+	W+/B−	D	W−/B+	W+/B−	D
Superior	933	711	222	951	780	171
Subordinate	890	638	252	880	848	32

$n = 11$ per condition.

bias. Positive difference scores revealed a bias in favor of white Americans. These difference scores, subsequently, were subjected to a 2 target race (white, black) × 2 role (superior, subordinate) Factorial ANOVA. The mean[4] differences representing automatic racial bias, as well as the mean latencies for each of the category pairing phases, are reported in Table 1.

Consistent with Lowery et al. (2001), participants revealed less bias if they anticipated interacting with a black individual, compared to a white individual [$F(1,40) = 14.82,\ p < .0004,\ r = .52$]. However, the main effect of target race was moderated by an interaction with situational role [$F(1,40) = 4.92,\ p < .05,\ r = .31$]. As predicted, participants exposed to a black subordinate (i.e., participants assigned to the superior role) revealed greater automatic racial bias ($M = 171$) than participants anticipating an interaction with a black superior ($M = 32$) [$t(40) = 1.98,\ p < .05,\ r = .38$]. By contrast, participants anticipating a same-race interaction did not differ in their degree of racial bias (respective Ms $= 222,\ 252$ for superior and subordinate roles, respectively) [$t(40) = 1.01,\ p < .32$]. Furthermore, participants in the subordinate role revealed greater bias if they anticipated a same-race interaction, rather than an interracial interaction [$t(40) = 4.21,\ p < .0001,\ r = .69$]; however, the bias of participants in the superior role did not differ as a function of the race of their interaction partner [$t(40) = 1.23, p < .23$]. These results suggest that the role white participants anticipated that they would hold for a dyadic interaction with a black individual influenced their degree of prejudice. That is, participants who anticipated holding a position of relatively low power for an interracial interaction generated attitudes that were less biased, compared to individuals in a powerful role, and those anticipating same-race interactions. Hence, both the demographics of the anticipated interaction, and participants' relative roles, impacted automatic racial attitudes.

Discussion

Prejudicial attitudes and stereotypical associations undoubtedly are systemic forces that are often activated and applied automatically, without the conscious awareness of the perpetrator (Bargh et al., 1992, 1996; Devine, 1989; Dovidio, Kawakami, Johnson, Johnson, & Howard, 1997; Fazio et al., 1995; Fazio et al., 1986; Greenwald & Banaji, 1995; see also Wilson, Lindsey, & Schooler, 2000). Automatic attitudes and prejudice directed toward social groups were originally thought immutable and resistant to the influence of contextual factors. The results of the present study add to a growing body of research suggesting otherwise—deeply rooted evaluations of social categories, it seems, can be influenced by features of the immediate social context (see Devine, 2001 for a summary). In particular, the present work highlights the important influence that situational roles can have on automatically activated attitudes (see also Richeson & Ambady, 2001). In the present study, participants' roles differentially impacted their automatic attitudes. We found that individuals holding a powerful position for an upcoming intergroup interaction were more biased than individuals holding a less powerful position. Situational power is known to influence the extent to which individuals engage in category-based information processing, such as stereotyping (Fiske, 1993). The findings of the present study extend this work to the domain of automatic attitudes.

The present findings also contribute to previous research examining the influence of exposure to black exemplars on whites' automatic racial bias (Dasgupta & Greenwald, 2001; Lowery et al., 2001; cf., Rudman et al., 2001[5]). Whereas Dasgupta and Greenwald (2001) find that individuals must be exposed to admired, atypical blacks in order to reduce bias, Lowery et al. (2001) argue that the actual presence of a black individual (whether atypical or not) may be sufficient. Although the presence of a black individual undoubtedly is a strong contextual cue, our results suggest the importance of that individual's typicality in the shaping of automatic racial attitudes. In Lowery et al.'s study, participants were not only presented with a black individual, but that black individual was in an atypical, authoritative role (i.e., the experimenter). A more compelling test of the "exposure alone" hypothesis requires

[4] For ease of interpretation, the means presented in the text and table have been re-transformed to millisecond units from the log-transformed means.

[5] Although they do not directly test the impact of black compared to white professors on white students' racial bias, Rudman et al. (2001) make the argument that exposure to a black professor without concurrent discussion of diversity will not reduce whites' racial bias.

exposure to a black individual holding a relatively low-status situational role. That is, the racial bias of whites exposed to a black individual holding a low-status role, and whites exposed to a white individual holding a low-status role, should differ. The attitudes of participants in the role of a "task superior" in the present study, in a sense, were exposed to a black individual in a low-status role (i.e., as their task subordinate). They did not, however, reveal less bias than participants exposed to a white subordinate. By contrast, participants anticipating an interracial interaction who were assigned to the subordinate role were exposed to a black individual in a high-status role. Drawing on the work of Dasgupta and Greenwald (2001), the black superior in the present study, and the black experimenter in Lowery et al.'s work, may have served as atypical exemplars for participants, leading to an attenuation of racial prejudice. Thus, the present findings provide a step toward disentangling the effects of exemplar typicality and exposure in the reduction of automatic prejudice.

Limitations

Several issues limit the generalizability of the present findings. First, the participants in the present work were all female. Had the gender of the target been different from the gender of the participants, a dimension of sociocultural status other than race (i.e., gender) may have influence the results. Consequently, we decided to control for gender in the present study. Future work should examine the automatic racial prejudice generated by participants of dyadic interactions that vary in both race and gender. Furthermore, the participants of the present study were students at an extremely competitive, private university. The racial attitudes held by these individuals may differ from other subsets of the general population of the country. The body of work revealing automatic attitude biases, however, has consistently found evidence for automatic racial prejudice, despite differences in the participant populations (Dasgupta et al., 2000; Fazio et al., 1995; Greenwald et al., 1998; Lowery et al., 2001; McConnell & Leibold, 2001; Ottaway et al., 2001; Rudman et al., 2001). Last, although the results of the present work provide strong evidence for the influence of situational factors on automatic attitude activation in anticipation of an interracial interaction, it remains unclear whether the actual presence of a black individual in a superior role would have influenced automatic attitudes in a similar manner and to the same degree as found in the present study. For instance, would the participants in the present work who revealed attitudes that were less biased in anticipation of an interracial interaction, also have revealed less bias after an interaction with a black superior, compared to participants in other roles? Future work should investigate the degree to which situational cues alter automatic attitudes both in anticipation of, as well as during and after, actual interracial dyadic interactions.

Implications and conclusion

Despite these limitations, the present study has broad implications for how we think about attitudes, and more importantly, attitude change. The present results suggest that attitudes are quite malleable, given the appropriate situation. In particular, this work suggests that reversing individuals' situational power during a dyadic interaction from what is typical in society may be one such situation. In the present research, the attitudes held by members of a dominant racial group were influenced by the reversal of power in the immediate situation from that maintained by society's *status quo*. Thus, this work suggests that changes in the attitudes of members of dominant social groups may be linked to their relative power during intergroup interactions—an often-ignored tenet of Allport's Contact Hypothesis (1954). Considered in this light, the findings inspire optimism regarding the reduction of prejudice and discrimination. Based on previous studies, prejudice reduction hinged on incidental exposure to famous, atypical exemplars of negatively stereotyped groups (Bodenhausen, Schwarz, Bless, & Wänke, 1995; Dasgupta & Greenwald, 2001; cf. Lowery et al., 2001), or more long-term, interpersonal contact with stigmatized individuals (Pettigrew, 1998; Rudman et al., 2001). Although such interventions are both hopeful and effective, both are somewhat limited insofar as most people are unlikely to incidentally be exposed to famous, atypical individuals or to voluntarily form friendships with members of stigmatized groups. Lack of exposure to atypical members of stigmatized groups, and resistance to interpersonal contact with them, may even be directly proportional to an individuals' degree of bias. But, the same individuals may incidentally and involuntarily be exposed to a member of a stigmatized group in an atypical or counter-stereotypical role, at work for instance. If one's black manager can have a similar impact on attitudes as Michael Jordan, then widespread prejudice reduction may be attainable.

Considered before the backdrop of current situational power hierarchies in workplaces, however, the present findings are fairly pessimistic. That is, given the extent to which situational power is correlated with sociocultural group status (Sidanius & Pratto, 1993), the present findings suggest that the most likely situational power arrangement for members of stigmatized and nonstigmatized groups during intergroup interactions serves to reinforce racial prejudice. Similar to the manner in which the cognitive biases of powerful people perpetuate group status differences during intergroup interactions, the present findings suggest that the attitudes automatically generated by powerful people overwhelmingly support social stratification as well.

Clearly, this work has important implications for diversity and Affirmative Action initiatives. A thoughtful examination of how, when, and to what end situational factors influence automatic attitudes may be vital for the attenuation of prejudice and intergroup conflict.

Acknowledgments

This research was supported by a National Institute of Mental Health Grant (#MH12212-01A1), a PECASE Award from the National Science Foundation (BCS-9733706), and a Burke Award from Dartmouth College. The authors would like to thank Anthony Greenwald for the Implicit Association Test software program, as well as Heather Gray and two anonymous reviewers for their comments on previous versions of this manuscript. Special thanks are also offered to Lauren Grabowski, Inga Budish, and Natalie Hornak for their assistance with data collection and stimulus preparation.

References

Allport, G. (1954). *The nature of prejudice*. Reading, MA: Addison-Wesley.

Bargh, J. A., Chaiken, S., Govender, R., & Pratto, F. (1992). The generality of the automatic attitude activation effect. *Journal of Personality and Social Psychology, 62*, 893–912.

Bargh, J. A., Chaiken, S., Raymond, P., & Hymes, C. (1996). The automatic evaluations effect: Unconditional automatic attitude activation with a pronunciation task. *Journal of Experimental Social Psychology, 32*, 104–128.

Blair, I. V., Ma, J. E., & Lenton, A. P. (2001). Imagining stereotypes away: The moderation of implicit stereotypes through mental imagery. *Journal of Personality and Social Psychology, 81*, 828–841.

Bodenhausen, G. V., Schwarz, N., Bless, H., & Wänke, M. (1995). Effects of atypical exemplars on racial beliefs: Enlightened racism or generalized appraisals? *Journal of Experimental Social Psychology, 31*, 48–63.

Dasgupta, N., McGhee, D. E., Greenwald, A. G., & Banaji, M. R. (2000). Automatic preference for White Americans: Eliminating the familiarity explanation. *Journal of Experimental Social Psychology, 33*, 316–328.

Dasgupta, N., & Greenwald, A. G. (2001). On the malleability of automatic attitudes: Combating automatic prejudice with images of admired and disliked individuals. *Journal of Personality and Social Psychology, 81*, 800–814.

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

Devine, P. G. (2001). Implicit prejudice and stereotyping: How automatic are they? Introduction to the special section. *Journal of Personality and Social Psychology, 81*, 757–759.

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.

Erber, R., & Fiske, S. T. (1984). Outcome dependency and attention to inconsistent information. *Journal of Personality and Social Psychology, 47*, 709–726.

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.

Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology, 50*, 229–238.

Fiske, S. T. (1993). Controlling other people: The impact of power on stereotyping. *American Psychologist, 48*, 621–628.

Goodwin, S. A., Gubin, A., Fiske, S. T., & Yzerbyt, V. Y. (2000). Power can bias impression processes: Stereotyping subordinates by default and by design. *Group Processes & Intergroup Relations, 3*, 227–256.

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review, 102*, 4–27.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association task. *Journal of Personality and Social Psychology, 74*, 1464–1480.

Jost, J., & Banaji, M. R. (1994). The role of stereotyping in system-justification and the production of false consciousness. *British Journal of Social Psychology, 33*, 1–27.

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

McConnell, A. R., & Leibold, J. M. (2001). Relations between the Implicit Association Test, explicit racial attitudes, and discriminatory behavior. *Journal of Experimental Social Psychology, 37*, 435–442.

Operario, D., Goodwin, S. A., & Fiske, S. T. (1998). Power is everywhere: Social control and personal control both operate at stereotype activation, interpretation, and response. In R. S. Wyer (Ed.), *Advances in social cognition* (Vol. 11, pp. 163–176). Mahwah, NJ: Erlbaum.

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

Pettigrew, T. F. (1998). Intergroup contact theory. *Annual Review of Psychology, 49*, 65–85.

Purdue, C. W., & Gurtman, M. B. (1990). Evidence for automaticity in ageism. *Journal of Experimental Social Psychology, 26*, 199–216.

Raven, B. H. (1993). The bases of power: Origins and recent developments. *Journal of Social Issues, 49*, 227–251.

Richeson, J. A., & Ambady, N. (2001). Who's in charge: Effects of situational roles on automatic gender bias. *Sex Roles, 44*, 493–512.

Rudman, L. A., Ashmore, R. D., & Gary, M. L. (2001). "Unlearning" automatic biases: The malleability of implicit prejudice and stereotypes. *Journal of Personality and Social Psychology, 81*, 856–868.

Sidanius, J., & Pratto, F. (1993). The dynamics of social dominance and the inevitability of oppression. In P. Sniderman & P. E. Tetlock (Eds.), *Prejudice, politics, and race in American today* (pp. 173–211). Stanford, CA: Stanford University Press.

Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review, 107*, 101–126.