

Moral Credentials and the Expression of Prejudice

Benoît Monin and Dale T. Miller
Princeton University

Three experiments supported the hypothesis that people are more willing to express attitudes that could be viewed as prejudiced when their past behavior has established their credentials as nonprejudiced persons. In Study 1, participants given the opportunity to disagree with blatantly sexist statements were later more willing to favor a man for a stereotypically male job. In Study 2, participants who first had the opportunity to select a member of a stereotyped group (a woman or an African American) for a category-neutral job were more likely to reject a member of that group for a job stereotypically suited for majority members. In Study 3, participants who had established credentials as nonprejudiced persons revealed a greater willingness to express a politically incorrect opinion even when the audience was unaware of their credentials. The general conditions under which people feel licensed to act on illicit motives are discussed.

One constant source of social anxiety in mainstream American society is the fear of being labeled sexist or racist. These labels hang perilously over every social interaction, threatening to drop on the head of any person who takes a false step. So debilitating is the concern with appearing prejudiced that Crocker, Major, and Steele (1998) described it as “akin to the stereotype threat experienced by the stigmatized” (p. 540). The source of this unease and wariness can be traced to powerful and pervasive social norms against the expression of prejudice that have emerged in recent decades (Devine, 1989; Fiske, 1998; McConahay, 1986; Smith, 1985). Whether these norms have reduced the level of actual prejudice is a matter of continued debate, but few deny that they have reduced the level of expressed prejudice (Gilbert, 1951; Jones & Sigall, 1971; Karlins, Coffman, & Walters, 1969; D. Katz & Braly, 1935; McConahay, 1986; McConahay, Hardee, & Batts, 1981). Fear of immediate social disapprobation is not the only reason people resist expressing negative beliefs about out-groups. Many individuals have internalized antiprejudice norms and closely monitor themselves to avoid behaving in ways that might clash with their egalitarian values (Devine, Monteith, Zuwerink, &

Elliot, 1991; Dutton & Lake, 1973; Gaertner & Dovidio, 1986; I. Katz & Hass, 1988; Myrdal, 1944; Plant & Devine, 1998).

One manifestation of people’s anxiety about appearing prejudiced (to themselves or to others) is their hesitancy to act on nonegalitarian attitudes unless those attitudes can plausibly be attributed to something other than prejudice. A series of studies by Snyder and his colleagues (Snyder, Kleck, Strenta, & Mentzer, 1979) illustrates this phenomenon. In Snyder et al.’s paradigm (see also Batson, Flink, Schoenrade, Fultz, & Pych, 1986), participants are asked to evaluate a movie playing on one of two monitors separated by a partition. Already sitting in front of each of the monitors is a confederate, who is either “normal” or stigmatized (e.g., disabled, Black). The critical manipulation is whether both monitors show the same movie or two different ones. Although many participants presumably wish to avoid sitting with the stigmatized other, only those in the two-movie condition appear comfortable acting on this attitude. The reason for this is that the two-movie condition, unlike the one-movie condition, renders the motivation for the participant’s seating choice ambiguous: It could reflect either the illicit motive of wishing to avoid the stigmatized other or the more legitimate motive of genuinely preferring the movie being shown on the monitor watched by the normal other.

The desire to avoid being prejudiced does not manifest itself only in the hesitancy to express nonegalitarian attitudes. Sometimes it can actually lead to the granting of preferential treatment to minority members. This so-called reverse discrimination occurs when actors fear that treating a minority member as they would a majority member could be attributed to prejudice. Dutton’s (1971) investigation of dress code enforcement at upscale restaurants illustrates this phenomenon. When couples arrived at the restaurants with the man wearing attire that was slightly too casual (a turtleneck sweater) they were nearly twice as likely to be seated when they were Black as when they were White. Although the maître d’s had a perfectly valid reason to refuse admission to the Black couples (i.e., the absence of a tie), they were not secure enough to assume the risk that their refusal would be attributed to prejudice.

Benoît Monin and Dale T. Miller, Department of Psychology, Princeton University.

This research was conducted while Benoît Monin was a fellow of the Princeton University Woodrow Wilson Society of Scholars, and it constituted his master’s thesis under the supervision of Dale T. Miller. The research was supported by National Institute of Mental Health Grant MH44069. Portions of this article were presented at the annual meeting of the American Psychological Society, Washington, DC, May 1998, and at the annual meeting of the Midwestern Psychological Association, Chicago, April 1999. We are particularly indebted to Kevin Carlsmith, Robyn Leboeuf, Leif Nelson, Michael Norton, and Jacquie Vorauer for their insightful comments on a draft of this article. We thank Alison Baer for her assistance with the data collection in Study 1, Kimberlee Weaver and Michael Norton in Study 2, and Kim Kelly in Study 3.

Correspondence concerning this article should be addressed to Benoît Monin, Department of Psychology, Green Hall, Princeton University, Princeton, New Jersey 08544. Electronic mail may be sent to bmonin@princeton.edu.

Self-Licensing Through Moral Credentials

Bowing to antiprejudice impulses or political correctness pressures prevents some problems for actors, but it creates others. For one, the act of discriminating in favor of minority individuals, although it spares people the worry of being or appearing prejudiced, can leave them feeling bad about themselves—for having failed to uphold a rule they were responsible for enforcing, for having treated both minority and majority members unfairly, for having lacked the willpower to resist social pressure, and so on. For prejudiced people, merely treating minority members fairly comes at a cost because it clashes with their firmly held belief that minority members do not deserve the same treatment as majority members do.

Political correctness pressures and strong egalitarian values thus yield costs at the individual level as much as they yield benefits at the interpersonal and collective levels. From this perspective, it is of particular interest to determine when people will ever override such pressures and freely speak their mind. One critical factor in this regard, and the focus of the present research, is the actor's confidence that his or her politically incorrect behavior can be attributed to something other than prejudice. But how do people acquire such security? We propose that one means of doing so is by establishing *credentials* as nonprejudiced people. For example, the more a man has shown that he is not a sexist, the less he will fear that his current behavior might be attributed to sexism and the more comfortable he will be expressing a pro-male attitude. "Moral credentials" of this type are most available to individuals who customarily behave in a nonprejudiced fashion. However, even prejudiced individuals sometimes engage in ostensibly nonprejudiced behavior, if only out of deference to the antiprejudice norms described earlier, and, thus, they too can establish and be liberated by moral credentials.

Consequences of Feeling Prejudiced

Previous research has demonstrated that giving people bogus feedback about their level of prejudice affects their future behavior. Participants who are told that they are prejudiced generally try to disconfirm this negative label by being more sympathetic to members of the target category. For example, Dutton and Lake (1973; see also Dutton & Lennox, 1974) shook the confidence of students who rated themselves low on prejudice by giving them false biofeedback indicating that they were highly aroused after viewing slides of interracial couples. As participants left the laboratory with their payment, either a White or a Black panhandler approached them. Those participants who had been threatened by the racist feedback gave more money to the Black panhandler than did those who had not been threatened. Sherman and Gorkin (1980) went one step beyond this type of passive labeling and actually tricked egalitarian participants into revealing stereotypical thinking. Participants had to resolve the apparent inconsistency of a surgeon crying, "It's my son," when a boy who just lost his father in a car accident is brought in the operating room. Even though the participants had been selected to be highly egalitarian, they often failed to consider the most parsimonious explanation, that the surgeon is the boy's mother. As a result, they felt guilty and uncomfortable and bolstered their profeminist attitudes in later tasks, as if to make up for their temporary lapse. More recently,

authors have developed sophisticated questionnaire-format tools that guide participants in their self-examination and force them to recognize the gap between what they believe and what they preach in the domain of prejudice. The special focus of this latter research is the affective consequence of *self-discrepancy* manipulations (e.g., Devine et al., 1991; Plant & Devine, 1998).

The Present Research

In the present research we seek to show that people whose past behavior provides them with moral credentials in that particular domain are more willing to voice opinions that violate the dictates of political correctness in that same domain. The present research departs from previous research in two important ways. First, whereas the emphasis in prior work has been on the consequences of feeling prejudiced, we consider the opposite, the consequences of feeling nonprejudiced. Second, rather than focusing on those cases in which someone else provides people with evidence that they are unprejudiced (e.g., Dutton & Lake, 1973), we focus on those cases in which people, through their prior actions, provide themselves with such evidence.

Each of the present three studies tests the hypothesis that people who have previously expressed antiprejudiced attitudes are more likely to express their true attitudes in contexts in which there exists the potential for accusations of prejudice. Study 1 examines the consequences of disagreeing with blatantly sexist statements on voicing a sexist-sounding preference. Study 2 studies the impact of making a nonprejudiced hiring decision on voicing a prejudiced-sounding preference. Study 3 examines whether the ability of credentials to promote the honest expression of opinions requires that the audience for this expression be aware of the credentials.

Study 1

Consider an individual who must choose between a man and a woman for a position or an award. People in this circumstance frequently experience discomfort, the source of which is their knowledge that they will be vulnerable to charges of sexism should they select the male candidate and possibly the creeping suspicion that these charges might be warranted. Individuals in this circumstance who genuinely believe that the man is the superior candidate face an interesting dilemma: Should they act on their true preference, thereby risking feeling sexist, or should they provide the politically correct answer, thereby misrepresenting their true attitude and possibly making a suboptimal choice? In study 1 we presented participants with just this predicament and examined how their prior record in the domain in question affected their resolution of the predicament. We hypothesized that participants would be more likely to explicitly express an implicit preference for a man when their previous actions provided them with non-sexist credentials.

Method

Participants. A female experimenter approached 202 Princeton undergraduates on campus (115 men, 87 women) and asked them to answer an anonymous two-page survey. One man and one woman failed to complete the survey and were dropped from the analysis.

Design and procedure. Participants first indicated whether they considered each of five negative statements about women to be right or wrong. We manipulated the wording of these statements such that all five of them referred either to “most” or to “some” women. The exact wording of the statements in the two conditions was as follows:

- Some/Most women are better off at home taking care of the children.
- Men are more emotionally suited for politics than are some/most women.
- The best job for some/most women is something like cook, nurse, or teacher.
- Some/Most women need a man to protect them.
- Some/Most women are not really smart.

The “some” versus “most” manipulation was inspired by a study by Salancik and Conway (1975) that manipulated agreement to behavioral statements by phrasing them either with the qualifier “frequently” or with “on occasion.” Following Salancik and Conway’s logic, we expected that the “some” and the “most” items would yield different rates of disagreement. Specifically, we expected participants to answer *wrong* to more statements when the list used the “most” wording than when it used the “some” wording. Because they disagreed with more statements, the latter group of participants would presumably feel that they had stronger credentials as nonsexists and be correspondingly more willing to voice a politically incorrect preference.

Participants in the base-rate group did not receive any statements prior to completing the dependent measure. Study 1 thus included three conditions: the “most” wording condition ($n = 60$), the “some” wording condition ($n = 59$), both of which included the five initial statements to be rated, and the base-rate condition ($n = 81$), which did not.

Once they had registered their agreement or disagreement with the “most/some” statements, participants completed three filler items about the building industry (e.g., “Are you at all familiar with the building industry?”) and then read the following scenario:

Imagine that you are the manager of a small (45-person) cement manufacturing company based in New Jersey. Last year was a particularly good one, and after you invested in increasing the output capacity of your plant, you decide that it would be very fruitful if you could find clients in other states to increase your business. Because you cannot spend too much time away from the plant, you decide to appoint someone to go around to prospective clients and negotiate contracts. This is a highly specialized market, and the job will mostly consist in going from one building site to another, establishing contacts with foremen and building contractors. It is also a highly competitive market, so bargaining may at some points be harsh. Finally, it’s a very technical market, and a representative that did not exude confidence in their technical skills would not be taken seriously by potential clients. Realizing how useful such a help would be for you, you decide to give the person chosen one of the top-five salaries in your company. Do you feel that this job is better suited for one gender rather than the other?

The scenario intentionally focused on a stereotypically male job environment and used gender-biased language (e.g., “foremen”) to lead participants to believe that a man was most appropriate for the job. Participants answered the final question on a 7-point continuous scale ranging from *yes, much better for women* (−3) to *yes, much better for men* (3) with a midpoint labeled *no, I do not feel this way at all*.

Results

Manipulation check. To determine whether the phrasing manipulation had the expected effect on the rate of statement dis-

agreement, we submitted the number of the “some/most” statements that participants disagreed with to a Credentials \times Gender analysis of variance (ANOVA). One participant was dropped from this analysis because he agreed with four items and left the fifth item blank. The credentials manipulation had the anticipated effect on disagreement rates: Participants responded *wrong* to more statements in the “most” condition ($M = 4.15$) than in the “some” condition ($M = 2.49$), $F(1, 114) = 49.1$, $p < .001$. There was neither a significant gender main effect nor a Gender \times Credentials interaction on this measure.

Effect of credentials. We hypothesized that participants in the “most” condition would be more likely than would those in the “some” condition to deviate from the politically correct midpoint option and express a preference for hiring a man. To test this hypothesis, we analyzed participants’ preferences in a 2×3 factorial ANOVA, crossing gender with condition (most, some, and base rate). Consistent with predictions, this analysis yielded a main effect for credentials, $F(2, 194) = 4.4$, $p < .05$. As Figure 1 illustrates, the main effect for credentials reflects the fact that participants favored a man more in the “most” condition ($M = 4.8$) than they did in either the “some” or the base-rate condition ($M_s = 4.3$ and 4.5 , respectively).

We also found a main effect for gender, $F(1, 194) = 9.9$, $p < .005$, and a Gender \times Credentials interaction, $F(2, 194) = 3.7$, $p < .05$. This unpredicted interaction appears to be due to the fact that only men were influenced by the credentials manipulation. Among men, the “most” condition ($M = 5.1$) differed significantly both from the “some” condition ($M = 4.4$) and from the base-rate condition ($M = 4.6$), $t(54) = 3.3$, $p < .005$, and $t(86) = 2.7$, $p < .01$, respectively, whereas the latter two did not differ from each other, $t(82) = -1.3$, $p > .2$. Among women, however, there were no significant differences across the three groups ($M = 4.4$, 4.3 , and 4.4 , respectively), all $t_s < 1$. To avoid capitalizing on chance through multiple pairwise comparisons, we also conducted within each gender a Tukey honestly significant difference test, a post hoc analysis that considers all possible pairwise comparisons simultaneously while maintaining an overall alpha of .05. With this more conservative test, we found the same pattern as before: Only among men did the “most” condition differ from the other two conditions, which did not differ from each other.

Discussion

The results from Study 1 generally support the credentials hypothesis. Male participants who were provided with the opportunity to disagree with five blatantly sexist statements were subsequently more likely to indicate that a job was better suited for a man than were male participants who first responded to more ambiguous (and thus less easy to reject) statements about women. Presumably, the opportunity to disagree with sexist items left the former participants feeling they had demonstrated that they were not sexist, thereby liberating them to respond more honestly on the dependent measure.

The only unexpected finding in Study 1 was that the credentials effect only emerged among men. We included female as well as male participants in this study because we thought it naive to assume that women are free of the concern of being or appearing sexist. For one, targets of prejudice often internalize the prejudiced beliefs that others hold of them (Apfelbaum, 1979; Jost & Banaji,

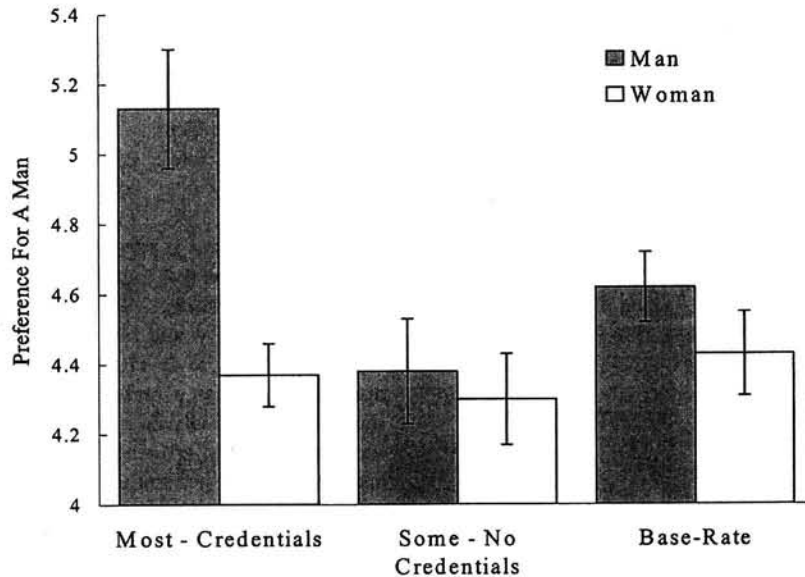


Figure 1. Mean preference for a man (\pm SE) by credentials and gender in Study 1.

1994; Triandis et al., 1982), which means that anti-Black racism can be found among African Americans, and sexism among women. Also, there is little evidence that women differ from men in their desire to avoid being or being seen as sexist—and one might even expect them to be more wary than men of being or appearing sexist.¹ For these reasons, we expected that credentials would operate similarly in men and women.

One possible explanation of the finding that nonsexist credentials did not affect women is that our manipulation of credentials had a different psychological effect on women. Specifically, it may have constituted a personal threat to women, one that made them defensively “freeze” on the egalitarian response options. Thus, the opportunity to disagree with items in the “most” condition (e.g., “Most women are not really smart”), although it left men feeling good about themselves, might have left women feeling insulted and demeaned. Indeed, some female participants actually expressed annoyance at the items after filling them out.

One possible alternative explanation for the observed effects of credentials deserves consideration. We speak here of the impact of merely reading, as opposed to rating, the “most/some” items. Specifically, mere exposure to the more blatant sexist statements in the “most” condition may have disposed participants to give more sexist responses to the cement factory scenario, just as hearing someone condone prejudice can lead individuals to adopt more prejudiced positions (Blanchard, Crandall, Brigham, & Vaughn, 1994). To address this possibility, we asked a separate group of 39 male participants to complete a modified version of the procedure: Instead of judging whether each item was right or wrong, they merely rated its ambiguity. The objective was to make participants read the items carefully yet not give them the opportunity to establish credentials by rejecting any item. We combined the responses of these 39 men with those of the 56 men who were in the “most” and “some” standard conditions discussed previously and submitted them to a Credentials \times Type of Rating ANOVA. Neither main effect was significant ($F_s < 1$), but the interaction was highly significant, $F(1, 91) = 13.0, p < .001$. Consistent with

the credentials hypothesis, reading the statements without responding to them had no effect on participants’ willingness to give sexist responses.

Study 2

Study 2 provides a further test of the hypothesis that establishing credentials as a nonprejudiced person increases a person’s willingness to express a prejudiced attitude. Study 2 uses the same dependent measure as Study 1 but a different manipulation of credentials. We changed the way nonprejudiced credentials were established for two reasons. First, in Study 2 we sought to demonstrate the broad applicability of the concept of moral credentials by showing that they operate in the domain of racism as well as sexism. Second, in Study 2 we sought to manipulate credentials in a manner that would not threaten female participants and therefore would yield comparable effects across the two genders. The manipulation of credentials in Study 2 involved either providing or not providing participants with the opportunity to select a woman or a minority member in an initial job selection task. We hypothesized that the act of recommending the hiring of a woman [an African American] on one job selection task would increase participants’ willingness to express the belief that a man [a White] was more suited for a position described on a second job selection task.

¹ In their investigations of the motivations to respond without sexism, Klonis and Devine (2000) include only male participants, thus leaving this question open. To address it, we gave 55 participants (30 women) drawn from the same population as were those in the studies below Klonis and Devine’s Internal and External Motivation to Respond Without Sexism Scales (IMS-S and EMS-S, respectively). Women did not show less concern about others seeing them as sexist (EMS-S), $t(53) = 0.52, ns$, but did show more internalized motivation not to actually be sexist (IMS-S), $t(53) = -3.3, p < .005$. In other words, women did not feel licensed vis-à-vis others merely because of their social category and, if anything, showed more personal concern about acting in a sexist way.

Method

Participants. One hundred thirty-two Princeton undergraduates (50 men, 82 women) took part in this study for payment or for experimental credit. Half of these participants were administered the sexism version of the procedure (35 credentials, 31 controls), and half were administered the racism version (36 credentials, 30 controls).²

Procedure and design. In the laboratory, participants made two recruitment-related decisions. The first of these required participants to indicate which of five applicants they would choose for a starting position in a large consulting firm. Each candidate was briefly described by means of a picture, name, college, grade point average (GPA), and major. In all conditions, the fourth applicant was designed to be the most attractive: He or she had graduated from a prestigious institution, had majored in economics, and had the highest GPA. The manipulated variable was the gender [ethnicity] of that star applicant. In the nonsexist credentials condition, the star applicant was a White woman; in the nonracist credentials condition, the applicant was an African American man; in the no-credentials control condition, the applicant was a White man. We manipulated the target's gender or race by changing his or her photograph (and first name in the sexism version). All four other applicants were invariably White men. Once participants had completed the first recruitment task, they completed three filler items and then responded to a dilemma designed to elicit the expression of prejudice. In the sexism condition, we used the cement factory dilemma used in Study 1. For the racism condition, we designed an analogue set in a racist police force:

Imagine that you are the police chief of a small town in a rural area of the U.S. Historically the population of the town has been exclusively White, and attitudes towards other ethnicities tend to be unfavorable. As much as you regret it, you know this is especially the case within your unit. You couldn't help overhearing racist jokes coming from people you otherwise consider excellent officers. In fact a couple of years ago an African-American patrolman joined your unit, and within a year he quit, complaining about hostile working conditions. You are doing what you can to change attitudes, but your main objective is that the police force should do its job, and so far it has been rather effective so you do not want to provoke any major unrest within the ranks. The time has come to recruit a new officer. As a general rule, officers need to be responsible and trustworthy, show quick intelligence enabling them to make split-second decisions in crisis situations. Recent scandals have also highlighted the need for a high level of integrity, resistance to corruption, mild manners and a calm temper. You have just received applications from the new graduates of the local Police Academy. You wonder whether ethnicity should be a factor in your choice. Do you feel that this specific position (described above) is better suited for any one ethnicity?

Like the scenario used in Study 1, this scenario was designed to make the working environment seem hostile to the stereotyped group, so that participants would have a legitimate reason to favor a White applicant. Participants answered the final question on a 7-point continuous scale ranging from *yes, much better for a Black* (-3) to *yes, much better for a White* (3), with a midpoint labeled *no, I do not feel this way at all*.

Results

Manipulation check. The manipulation proved effective. Over 83% of the participants chose Edwards, the star applicant, and this proportion did not differ by credentials condition, $\chi^2(1, N = 132) = 1.8, ns$, nor by prejudice type, $\chi^2(1, N = 132) = 0.9, ns$. Although it would have been desirable for all participants to have chosen Edwards, the fact that 17% chose a candidate other than Edwards only diminishes the chances of finding the predicted effect. The reason for this is that by choosing a candidate other

than the star applicant in the credentials condition, participants inevitably chose a White man, as they always did in the no-credentials condition, whether that White man was Edwards or someone else. Nevertheless, to avoid a self-selection bias, we first analyzed the results including all 132 participants and then repeated the analysis dropping the troublesome participants.

Effect of credentials. We conducted a 2 (credentials) \times 2 (type of prejudice) \times 2 (gender) ANOVA on the responses to the job suitability question. The only significant effect to emerge was the predicted main effect for credentials, $F(1, 124) = 6.3, p < .05$, all other F s *ns*. As Figure 2 shows, participants were more likely to favor a White man in the credentials condition ($M = 4.8$) than in the noncredentials condition ($M = 4.4$). The fact that credentials did not interact with type of prejudice, $F(1, 124) = 0.2, ns$, indicates that the manipulation had comparable effects on the racism and sexism measures.

We indicated earlier that a small minority of participants (22 of 132) failed to choose the target applicant. When we reanalyzed the data omitting the responses of these participants, a main effect for credentials emerged once again, $F(1, 102) = 9.6, p < .005$, all other F s *ns*. As expected, participants favored a White man more in the credentials condition ($M = 4.9$) than in the no-credentials condition ($M = 4.3$).

Discussion

Study 2 replicates the effect of moral credentials observed in Study 1. Participants who established nonprejudiced credentials by selecting a woman [an African American] in the first recruitment task were more willing to indicate that a man [a White] was better suited for a second job than were participants who did not have the opportunity to recommend a woman [an African American] applicant in the first task. This study resembles a real workplace situation, in which decision makers are presented with series of choices involving female or African American coworkers. It suggests that a decision that favors one minority member (even if it is totally deserved) is sufficient to liberate people to act on an attitude (often based mainly on prejudice) that is detrimental for other minority members.

Unlike in Study 1, the manipulation of credentials used in Study 2 affected men and women equally. Whereas only men showed the effect of nonsexist credentials in Study 1, both men and women showed the effect in Study 2. In addition to the finding that gender did not interact significantly with either credentials or type of prejudice, an analysis that considered only the responses of participants who completed the sexism part of the design ($n = 66$) revealed no significant two-way interaction between credentials and gender, $F < 1$. These results support the claim that the Gender \times Credentials interaction observed in Study 1 was due to the way credentials were operationalized in that study and not to

² Whereas both the sexism and the racism version include female participants, we excluded African American participants from the racism version because of the limited number of African Americans in our participant population. We expect, however, that nonracist credentials would operate similarly among African Americans, for the same reasons that we expect nonsexist credentials to influence women.




Figure 2. Mean preference for a White person [a man] (\pm SE) by prejudice type and credentials in Study 2.

the fact that men have a greater susceptibility to nonsexist credentials than do women.³

Study 3

One interesting question about credentials that has not been addressed thus far is whether they function solely to reduce people's fear of appearing prejudiced or also function to allay personal concerns about actually being prejudiced. In Studies 1 and 2 participants' responses were always given in writing, and participants were assured that their answers would remain anonymous and confidential. When researchers manipulate the public/private nature of a response, they often follow such a procedure in creating the private condition, asking participants to read their answers aloud to the experimenter in the public condition (e.g., Plant & Devine, 1998). Given that the effects reported so far were all obtained in such private conditions, it is tempting to argue that credentials at least partly serve to allay people's concerns about being prejudiced, not simply their concerns about appearing prejudiced to others. However, one could still argue that participants' greater willingness to express politically incorrect beliefs following the acquisition of credentials was entirely due to their diminished fear of appearing prejudiced to the experimenter, who was aware of both sets of responses. To rule out this possibility, in Study 3 we used different experimenters in the two phases of the experiment. The use of different experimenters to administer the manipulation and collect the dependent measure is a standard technique for determining whether an effect is based solely on impression management (e.g., Dutton & Lake, 1973; Freedman & Fraser, 1966, Study 2). If the effect vanishes when two experimenters are involved (e.g., Rosenberg, 1965), then it can be attributed to something akin to evaluation apprehension. Otherwise, some other mechanism has to be assumed.

The logic behind the design of Study 3 is as follows. To the extent that credentials operate at the intrapersonal level, it should not matter whether the audiences to the two sets of responses are the same or different. However, to the extent that credentials

operate at the interpersonal level, they should only be effective when the audience to the participant's behavior in the second context is also aware of his or her prior behavior. Stated more formally, a public image model of credentials predicts that establishing credentials will only license people to give politically incorrect responses under same-audience conditions, whereas a self-image model of credentials predicts that credentials will license people under both same-audience and different-audience conditions.

Method

Participants. Thirty-five non-Black Princeton undergraduates took part in the study for partial course credit. Four participants had to be dropped because they had already taken a version of the task, and 1 participant because he was suspicious. The 30 remaining participants (17 men, 13 women) were assigned randomly to three conditions.

Procedure. At the end of an unrelated experimental session, the experimenter presented participants with a questionnaire and set of instructions that closely paralleled those used in the racism condition of Study 2. The one deviation from the prior procedure was the audience manipulation. In the *credentials-one audience* condition, participants completed both the nonracist credentials version of the consulting firm selection task and the police force selection task for the same experimenter, as did participants in Study 2. In the *no-credentials* condition, participants completed the control version of the consulting firm selection task and then the police force selection task, again for the same experimenter. In the *credentials-two audiences* condition, participants first completed the nonracist version of the credentials consulting firm task for one experimenter, and then a second experimenter knocked on the door, asking whether there was time left to do his questionnaire. After the first experimenter checked her watch, she agreed to let the second experimenter hand out his task, which was the police force scenario. Participants in this condition were thus led to believe that the person reading their answer to the police force scenario would not

³ Because we did not include African American participants in the racism half of the design, we were unable to determine whether Blacks would show a licensing effect in a racially charged scenario.

be aware of their answer to the consulting firm task. Both experimenters were White. Afterward, participants completed a questionnaire that probed them for suspicion and then were fully debriefed.

Results

Manipulation checks. The suspicion questionnaire asked participants to indicate which experimenter had administered which part of the experiment to them and which experimenter was going to have access to each of their two sets of responses. For each of these four questions, participants could circle the name of the first experimenter, the name of the second experimenter, or *other*, and specify. The questionnaire thus was a fairly conservative test of suspicion, as it explicitly suggested to participants the possibility that someone other than the person giving them each part might be reading it. Yet only 1 participant in the two-audience condition (who was dropped from the analysis) suspected that both sets of responses might be seen by the same experimenter. Three participants in the one-audience condition supposed that both experimenters would be reading both tasks, but these participants' data were retained, as this did not represent a manipulation failure. We also checked whether participants actually chose the star candidate, John, on the first selection task. All but one participant picked John. We kept this participant in the analysis to avoid self-selection biases.

Expressed prejudice. To begin with, we checked to see whether the findings of Study 3 replicated those of previous studies. We did this by comparing the credentials and no-credentials conditions under the same-audience condition. The earlier findings were replicated in that participants who were given the opportunity to establish nonracist credentials favored a White more than did those who were not given the opportunity to establish credentials ($M = 4.9$ vs. 4.2), $t(19) = 2.2$, $p < .05$.

Second, we sought to see whether the credentials effect occurred even when the audiences for the two tasks were different. We did this by conducting two nonorthogonal contrasts (Rosenthal &

Rosnow, 1991). Given the credentials–one audience, no-credentials, and credentials–two audiences array, the $(1, -2, 1)$ contrast tests the prediction that the nature of the audience does not matter, whereas the $(2, -1, -1)$ contrast tests the prediction that the credentials effect does not occur when different audiences are involved. The first contrast was significant, $t(27) = 2.2$, $p < .05$, whereas the second contrast was not, $t(27) = 1.4$, *ns*. Indeed, as Figure 3 shows, the two-audience condition yielded responses that were virtually identical ($M = 4.8$) to those yielded by the one-audience condition ($M = 4.9$).

Discussion

Study 3 replicates the credentials effect demonstrated in the previous studies. After recruiting a Black applicant in one job selection task, participants were more likely to express a preference for a White in a second, unrelated job selection task. More important, Study 3 addresses a central question regarding the psychological processes underlying the moral credentials effect. Namely, do credentials allay the fear of being seen as prejudiced or the fear of seeing oneself as prejudiced? In Study 3 we tackled this question directly by including a condition in which the audience in front of whom the participants made their second choice was not aware of the previous credentials-establishing behavior. The results in this private condition parallel those in the standard credentials condition, in which others knew of the preexisting credentials. This finding suggests that it is not critical that others know of one's credentials for them to have a licensing effect. In other words, moral credentials do not serve solely to make one appear less prejudiced to others; they also serve, at least partially, to reaffirm one's self-image as a nonprejudiced person.

General Discussion

According to most contemporary theories of sexism (Glick & Fiske, 1996; Swim, Aikin, Hall, & Hunter, 1995; Tougas, Brown,




Figure 3. Mean preference ($\pm SE$) for a White person by condition in Study 3. Cred. = credentials.

Beaton, & Joly, 1995) and racism (Gaertner & Dovidio, 1986; I. Katz & Hass, 1988; McConahay, 1986; Pettigrew & Meertens, 1995; Sears, 1988), people struggle to avoid behaving in ways that reflect prejudice. One important determinant of people's willingness to express politically incorrect attitudes is their confidence that they are unprejudiced. Although confidence of this type may derive from a variety of sources, an especially important source is people's past behavior—their track record, so to speak. The more confident people are that their past behavior reveals a lack of prejudice, the less they will worry that their future behavior is, or can be construed as, prejudiced and the more willing they will be to provide politically incorrect responses. Thus, the stronger a person's nonsexist or nonracist credentials, the less worried he or she will be that a given action would constitute (or would be seen as constituting) evidence of a sexist or racist disposition. By fostering self-image security, therefore, the establishment of moral credentials emboldens the actor to respond honestly in circumstances in which political correctness pressure militates against honest expression. Credentials, it is important to note, do not leave the actor feeling more motivated to respond honestly; they simply leave him or her feeling more comfortable doing so. In the language of approach-avoidance theory (N. E. Miller, 1944) moral credentials serve to reduce the avoidance tendencies associated with honest expression, not to increase the approach tendencies associated with it. The results of the present three studies support our analysis. Participants who established nonsexist credentials, either by dismissing a series of blatantly sexist statements (Study 1) or by selecting a female or African American candidate in an initial job recruitment task (Studies 2 and 3), were subsequently more likely to express prejudiced-sounding attitudes.

The Symbolic Nature of Credentials

The present findings demonstrate not only that credentials can increase people's willingness to express a prejudiced attitude but that they can do so even when the behavior on which these credentials are based cannot reasonably be viewed as highly diagnostic of nonprejudice. This symbolic nature of credentials is especially clear in Studies 2 and 3. For instance, selecting the woman in the initial recruitment task cannot reasonably be viewed as diagnostic of one's attitudes toward women because the fact that she is so much better qualified than the male applicants means that rejecting her would be self-defeating. Yet participants who chose her acted as though they had made a clear statement about their lack of sexism, as evidenced by their subsequent greater willingness to say that the construction job was better suited for a man.

Of course, even if the nonsexist credentials had been based on more diagnostic evidence, the legitimacy of the use to which they were put could be questioned. Why should a single act of goodwill toward one woman in one employment context justify a person adopting a position that discriminates against all women in another employment context? Viewed in that light, the symbolic role played by credentials in Studies 2 and 3 is reminiscent of the practice of tokenism (Kanter, 1977), a phenomenon that occurs when organizations advertise the few minority members they have hired or promoted in the apparent belief that this proves that they are egalitarian. In both contexts, decision makers seem disposed to treat what is at most a molehill's worth of goodwill as though it demonstrates a mountain's worth of virtue.

Alternative Interpretations

Some of the credentials findings are reminiscent of effects found in stereotype suppression studies. These studies find that instructing people to suppress a stereotype during an initial task results in the stereotype having a stronger impact on a subsequent task in which no suppression instructions are given (Macrae, Bodenhausen, Milne, & Jetten, 1994). Might such a "rebound" process have been at work in the present studies? For example, is it possible that participants in the "most" condition in Study 1 had to suppress their stereotypes about women to disagree with items such as "most women are not really smart" and, having done this, found their stereotypes bouncing back when they responded to the job suitability question? Although it is possible, we think it is unlikely for two reasons. First, it is difficult to see why one would need to suppress stereotypes to disagree with blatantly sexist items. Participants in stereotype suppression studies typically describe a target person's day while suppressing the stereotype attached to the target's social group. Their task is to scan the ideas provided by their imagination (most of which are bound to be stereotype congruent) and to reject any that are stereotypical. The mental control required in this, which presumably is critical to the subsequent rebound effect, is clearly absent in the present tasks. Second, rebound effects are typically found on unobtrusive measures (e.g., lexical decision task, seating choice), as it is supposedly only when participants stop exercising mental control that so-called ironic processes occur (Wegner, 1994). In our studies, the dependent variable was hardly unobtrusive. Participants were fully aware that the cement factory recruitment task involved stereotypes, and this is precisely why the task was suited for our purposes. Consequently, it is difficult to imagine a participant feeling a need to suppress in the "most/some" task and then suddenly deciding to let go when he or she comes to the cement factory task two pages later. Finally, although the rebound model might be stretched to fit the data of Study 1, it cannot plausibly account for the results of Studies 2 and 3. No stereotype suppression is required for participants to select the much more qualified Black or female candidate for the job recruitment task. Still, it is possible that there are some situations in which rebound effects and credentials work jointly. Indeed, the "day in the life" manipulation so often used in rebound research may, with the right choice of dependent variable, be shown to have licensing effects.

One might also argue that the present findings merely reflect participants' desire to balance gender and race preferences across tasks. Thus, after recruiting a woman for one job, participants may feel an obligation to recommend a man for a subsequent job. This interpretation also seems unpersuasive, however. First, it is unclear what the unit in need of balance would be. Why balance disagreeing with five sexist items with favoring a man in Study 1? Why not seek balance within the items by disagreeing with some and agreeing with others? And why does selecting a Black man who is the best qualified applicant for the consulting position need to be counterbalanced by denying the suitability of African American applicants for a police force position? Second and more important, the balancing interpretation cannot explain the results of Study 2. If participants were merely balancing their choice of a woman in the consulting recruitment task when they favored a man in the cement factory recruitment task, then participants in the control

condition, who also pick a man in the consulting task, should also have favored a woman in the cement factory task. But they did not.

The Interaction of Credentials and Prejudice

Earlier, we proposed that moral credentials should affect high- and low-prejudice individuals similarly. This does not mean, however, that the processes underlying the phenomenon are the same in the two groups. First, the two groups establish credentials differently: Whereas low-prejudice individuals establish (genuine) credentials by behaving in line with their deeply held attitudes, high-prejudice individuals establish (bogus) credentials by merely complying with antiprejudice norms. Second, high- and low-prejudice individuals can be expected to use credentials for different purposes. High-prejudice individuals tend to use them to induce charitable interpretations from others for their prejudiced opinions and actions, whereas low-prejudice individuals tend to use them to avoid being paralyzed by a constant fear of saying or doing the wrong thing. These distinctions highlight the ambivalent nature of credentials: On the one hand, they enable people of good will to speak their mind without fear of being misinterpreted; on the other, they let bigots act on their prejudice in the hope that they can get away with it.

The preceding analysis suggests that the most important moderator of the credentials effect may not be individuals' level of prejudice but rather their specific motivation for acting in a non-prejudiced manner. People differ in their motivations for responding without prejudice (Dunton & Fazio, 1997): Some people have internal motivations, some people have external motivations, some people have both, and some people have neither (Plant & Devine, 1998). To investigate the potential moderating influence of these variables, we administered the racism condition of Study 2 to 191 participants who had been pretested at the beginning of the semester on Plant and Devine's motivation scales. We then performed a regression analysis on the dependent measure (suitability), including credentials, IMS, EMS, and all possible interactions by computing the product of the centered variables (Aiken & West, 1991). As predicted, we once again found a strong main effect for credentials, $t(183) = -3.8, p < .0005$. We also found a main effect for IMS, $t(183) = -2.9, p < .005$, and a marginal effect for EMS, $t(183) = -1.8, p < .08$. It is not surprising that those participants who were most strongly motivated to respond without prejudice gave the most egalitarian answers. It is interesting to note that neither scale interacted with credentials (both $ps > .6$), nor was there a significant three-way interaction ($p > .35$). In fact, participants falling in all of the four quadrants of the two-dimensional space defined by the IMS and EMS Scales showed equal sensitivity to credentials. These data suggest that, at least within the range of individual differences available in a college population, moral credentials affect people equally, irrespective of their motivation for responding without prejudice.

The Liberating Versus Constraining Effects of Past Behavior

Readers no doubt have noted that a number of social psychological theories yield predictions that are diametrically at odds with the present results. These *consistency theories* assume that action, rather than liberating people to act differently in the future,

commits people to acting similarly. How can the claim that past behavior is liberating be reconciled with the (well-documented) claim that past behavior is constraining?

To understand this tension better, consider a commonly evoked example of the self-committing impact of prior behavior—the so-called *foot-in-the-door* effect. The typical finding in a foot-in-the-door study (e.g., Freedman & Fraser, 1966) is that participants who voluntarily comply with a small request (e.g., answering a brief telephone survey about household products) are subsequently more willing to comply with a second, much larger request (e.g., consent to a 2-hr intrusive home inspection by a group of five men). The foot-in-the-door effect appears to depend on two steps (DeJong, 1979). First, by observing their own behavior in the initial, small compliance situation, people adopt an image of themselves as the kind of person who is generally cooperative when approached with that kind of request. Second, having made that self-attribution and being confronted with the more burdensome request, people feel compelled to respond in ways that maintain their self-image.

In both the foot-in-the-door and the present moral credentials paradigms, participants respond sequentially to two situations that assess the same motivation. Despite this similarity, the observed consequences seem to be exactly opposite. In the foot-in-the-door paradigm the consequence of being presented with the first situation is an increased likelihood of responding similarly in the second situation; in the present paradigm the consequence of being presented with the first situation is a decreased likelihood of responding similarly in the second situation. If the conflicting results were not puzzling enough, the accounts given for them are actually quite similar: Both focus on the impact that past behavior has on one's identity, in particular how it serves to strengthen a particular image one holds of oneself or that one hopes to be displaying to the world (e.g., being helpful, being nonsexist).

When does projecting a particular self-image commit one to act similarly in the future, and when does it liberate one to act differently? A definitive answer to this question must await subsequent research, but we can offer a few thoughts at this time. One critical distinction between the two paradigms is the relation between the first and second response in each. In the foot-in-the-door paradigm, people appear compelled to respond in the second situation as they did in the first because of their fear that to do otherwise would result in an unflattering redefinition of their first act. If one agrees to help a person in fair weather (i.e., when the request is small), refusing to do so in foul weather (i.e., when the request is large) puts one at risk of seeing oneself and being seen by others as a fair weather friend. In the case of the present paradigm the situation is the reverse. Its psychology depends not on the power that a particular action (noncompliance) in the second situation has to redefine the meaning of the previous action (compliance) but on the power that the first action has to redefine the meaning of the action taken in the second situation. Performing a nonsexist act does not mean that one is now able to perform a sexist act with impunity, only that a subsequent ambiguous act on one's part is less likely to be interpreted as sexist. In effect, the moral credentials effect relies on an intuitive theory of self-consistency: the belief held by actors that they are consistent and thus that any ambiguous future behavior of theirs can be considered consistent with their past behavior.

One implication of this interpretation is that changing the ambiguity of the second measure in the two paradigms should change the pattern of responses. In foot-in-the-door paradigms, if the meaning of the second request was made more ambiguous—for instance, if it aided a cause of dubious merit—compliance with an earlier request might lower (rather than raise) compliance with the second one. The reason for this is that the first act would provide the actor with credentials as a helpful person, thereby directing attribution for the subsequent act of noncompliance toward the peculiarities of the situation and away from the fickleness of the actor. Conversely, the manipulations presented in this article might actually have had a committing effect had the second request been more clearly about prejudice. The more unambiguous the second measure of sexism is, the more constraining and the less liberating is the initial nonsexist action.

Beyond Prejudice

In the complex world of social interaction, seldom is a situation so clear cut that no ambiguity remains as to the causes of one's behavior. Credentials can turn ambiguity to one's advantage by discrediting the least flattering of the plausible motivations and, thus, bringing the most legitimate ones to the fore. We focus on the expression of prejudice because it is the most studied of illicit motivations in current psychological research, but we could just as well have used any of the many motivations that most of us would rather not be or be seen to be guided by. In all of these domains we would expect moral credentials to play the same role, liberating people to engage in ambiguous behavior because of a diminished fear of seeing themselves or being seen by others in a negative light.

Much of the work on the concealment and expression of prejudice could productively be generalized to the concealment and expression of all illicit motives. Snyder et al. (1979) explicitly used prejudice merely as an instance of "motives that people wish to conceal" (p. 2297). Their ultimate aim was to design a general "motive detection strategy" (p. 2305). In fact, Bernstein, Stephenson, Snyder, and Wicklund (1983) used the exact same two-movie paradigm to study a very different motivation, replacing the stigmatized confederate with an attractive woman. They assumed that their male participants would be reluctant to admit wanting to sit next to the confederate just because of her good looks, and, indeed, they observed the exact opposite pattern of what Snyder et al. (1979) had observed with a stigmatized confederate. For instance, in their first study, 75% of participants sat away from the attractive woman when the two movies were the same and thus did not provide them with an excuse, whereas 75% sat next to her when the movies differed.

Future Directions

Past behavior is not the only means by which a person's moral credentials are established. People can acquire moral credentials through their group membership, through what others have done to them (e.g., few would reproach a medical malpractice victim for holding negative stereotypes about doctors), or simply by association (e.g., if one's mother was a civil rights activist). Acquiring credentials through the labeling process seems a particularly fruitful avenue for future research. A person who has been labeled

(credentialled) as a nonsexist may feel liberated from antiprejudice norms to an extent not experienced by someone not so labeled. Demonstrating such an effect would be especially interesting because previous research (e.g., Kraut, 1973; R. L. Miller, Brickman, & Bolen, 1975) has demonstrated that labels can constrain people. For example, labeling someone as a helpful person has been found to lead that person to be more helpful in the future (Miller et al., 1975). Resolving the question of when labeling is constraining and when it is liberating, like that of when prior behavior is constraining and when it is liberating, could provide valuable social psychological insights. There already exists some research suggesting that a simple consistency model cannot account for labeling effects in the context of prejudice: For instance, giving people false feedback indicating that they are prejudiced tends to motivate them to contradict that label, not to confirm it (Dutton & Lennox, 1974; Sherman & Gorkin, 1980). Other studies show that people can react against even a positive label and behave in precisely the opposite way from that implied by the label (Brehm & Brehm, 1981; Snyder & Wicklund, 1981). All these examples suggest that there is much to be learned about the differences between circumstances that liberate or empower people and those that constrain them.

References

- Aiken, L. S., & West, S. G. (1991). *Multiple regression: Testing and interpreting interactions*. Newbury Park, CA: Sage.
- Apfelbaum, E. (1979). Relations of domination and movements for liberation: An analysis of power between groups. In W. G. Austin & S. Worchel (Eds.), *The social psychology of intergroup relations* (pp. 188–204). Monterey, CA: Brooks/Cole.
- Batson, C. D., Flink, C. H., Schoenrade, P. A., Fultz, J., & Pych, V. (1986). Religious orientation and overt versus covert racial prejudice. *Journal of Personality and Social Psychology*, *50*, 175–181.
- Bernstein, W. M., Stephenson, B. O., Snyder, M. L., & Wicklund, R. A. (1983). Causal ambiguity and heterosexual affiliation. *Journal of Experimental Social Psychology*, *19*, 78–92.
- Blanchard, F. A., Crandall, C. S., Brigham, J. C., & Vaughn, L. A. (1994). Condemning and condoning racism: A social context approach to interracial settings. *Journal of Applied Psychology*, *79*, 993–997.
- Brehm, S. S., & Brehm, J. W. (1981). *Psychological reactance: A theory of freedom and control*. New York: Academic Press.
- Crocker, J., Major, B., & Steele, C. (1998). Social stigma. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (Vol. 2, pp. 504–553). New York: McGraw-Hill.
- DeJong, W. (1979). An examination of self-perception mediation in the foot-in-the-door effect. *Journal of Personality and Social Psychology*, *37*, 2221–2239.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 5–18.
- Devine, P. G., Monteith, M. J., Zuwerink, J. R., & Elliot, A. J. (1991). Prejudice with and without compunction. *Journal of Personality and Social Psychology*, *60*, 817–830.
- 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.
- Dutton, D. G. (1971). Reactions of restaurateurs to Blacks and Whites violating restaurant dress requirements. *Canadian Journal of Behavioural Science*, *3*, 298–302.
- Dutton, D. G., & Lake, R. A. (1973). Threat of own prejudice and reverse discrimination in interracial situations. *Journal of Personality and Social Psychology*, *28*, 94–100.

- Dutton, D. G., & Lennox, V. L. (1974). Effect of prior "token" compliance on subsequent interracial behavior. *Journal of Personality and Social Psychology, 29*, 65-71.
- Fiske, S. T. (1998). Stereotyping, prejudice, and discrimination. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (Vol. 2, pp. 357-412). New York: McGraw-Hill.
- Freedman, J. L., & Fraser, S. C. (1966). Compliance without pressure: The foot-in-the-door technique. *Journal of Personality and Social Psychology, 4*, 195-202.
- Gaertner, J. L., & Dovidio, J. F. (1986). The aversive form of racism. In J. F. Dovidio & J. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 61-89). New York: Academic Press.
- Gilbert, G. M. (1951). Stereotype persistence and change among college students. *Journal of Abnormal and Social Psychology, 46*, 245-254.
- Glick, P., & Fiske, S. T. (1996). The Ambivalent Sexism Inventory: Differentiating hostile and benevolent sexism. *Journal of Personality and Social Psychology, 70*, 491-512.
- Jones, E. E., & Sigall, H. (1971). The bogus pipeline: A new paradigm for measuring affect and attitude. *Psychological Bulletin, 76*, 349-364.
- Jost, J. T., & 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.
- Kanter, R. M. (1977). *Men and women of the corporation*. New York: Basic Books.
- Karlines, M., Coffman, T. L., & Walters, G. (1969). On the fading of social stereotypes: Studies in three generations of college students. *Journal of Personality and Social Psychology, 13*, 1-16.
- Katz, D., & Braly, K. (1935). Racial prejudice and racial stereotypes. *Journal of Abnormal and Social Psychology, 28*, 280-290.
- Katz, I., & Hass, R. G. (1988). Racial ambivalence and American value conflict: Correlational and priming studies of dual cognitive structures. *Journal of Personality and Social Psychology, 55*, 893-905.
- Klonis, S. C., & Devine, P. D. (2000, February). *Internal and external motivation to respond without sexism*. Poster session presented at the annual meeting of the Society for Personality and Social Psychology, Nashville, TN.
- Kraut, R. E. (1973). Effects of social labeling on giving to charity. *Journal of Experimental Social Psychology, 9*, 551-562.
- Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Jetten, J. (1994). Out of mind but back in sight: Stereotypes on the rebound. *Journal of Personality and Social Psychology, 67*, 808-817.
- McConahay, J. B. (1986). Modern racism, ambivalence, and the modern racism scale. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 91-125). New York: Academic Press.
- McConahay, J. B., Hardee, B. B., & Batts, V. (1981). Has racism declined in America? It depends on who is asking and what is asked. *Journal of Conflict Resolution, 25*, 563-579.
- Miller, N. E. (1944). Experimental studies in conflict. In J. M. Hunt (Ed.), *Personality and the behavior disorders* (pp. 431-465). New York: Ronald.
- Miller, R. L., Brickman, P., & Bolen, D. (1975). Attribution versus persuasion as a means for modifying behavior. *Journal of Personality and Social Psychology, 31*, 430-441.
- Myrdal, G. (1944). *An American dilemma: The Negro problem and modern democracy*. New York: Harper & Brothers.
- Pettigrew, T. F., & Meertens, R. W. (1995). Subtle and blatant prejudice in Western Europe. *European Journal of Social Psychology, 25*, 57-75.
- Plant, E. A., & Devine, P. G. (1998). Internal and external motivation to respond without prejudice. *Journal of Personality and Social Psychology, 75*, 811-832.
- Rosenberg, M. J. (1965). When dissonance fails: On eliminating evaluation apprehension from attitude measurement. *Journal of Personality and Social Psychology, 1*, 28-42.
- Rosenthal, R., & Rosnow, R. L. (1991). *Essentials of behavioral research: Methods and data analysis* (2nd ed.). New York: McGraw-Hill.
- Salancik, G. R., & Conway, M. (1975). Attitude inferences from salient and relevant cognitive content about behavior. *Journal of Personality and Social Psychology, 32*, 829-840.
- Sears, D. O. (1988). Symbolic racism. In P. A. Katz & D. A. Taylor (Eds.), *Eliminating racism: Profiles in controversy* (pp. 53-84). New York: Plenum Press.
- Sherman, S. J., & Gorkin, L. (1980). Attitudes bolstering when behavior is inconsistent with central attitudes. *Journal of Experimental Social Psychology, 16*, 388-403.
- Smith, A. W. (1985). Cohorts, education & the evolution of tolerance. *Social Science Research, 14*, 205-225.
- Snyder, M. L., Kleck, R. E., Strenta, A., & Mentzer, S. J. (1979). Avoidance of the handicapped: An attributional ambiguity analysis. *Journal of Personality and Social Psychology, 37*, 2297-2306.
- Snyder, M. L., & Wicklund, R. A. (1981). Attribute ambiguity. In J. H. Harvey, W. Ickes, & R. F. Kidd (Eds.), *New directions in attribution research* (Vol. 3, pp. 197-221). Hillsdale, NJ: Erlbaum.
- Swim, J. K., Aikin, K. J., Hall, W. S., & Hunter, B. A. (1995). Sexism and racism: Old-fashioned and modern prejudice. *Journal of Personality and Social Psychology, 68*, 199-214.
- Tougas, F., Brown, R., Beaton, A. M., & Joly, S. (1995). Neosexism: Plus ça change, plus c'est pareil. *Personality and Social Psychology Bulletin, 21*, 842-849.
- Triandis, H. C., Lisansky, J., Setiadi, B., Chang, B. H., Marin, G., & Betancourt, H. (1982). Stereotyping among Hispanics and Anglos: The uniformity, intensity, direction, and quality of auto- and hetero-stereotypes. *Journal of Cross-Cultural Psychology, 13*, 409-426.
- Wegner, D. M. (1994). Ironic processes of mental control. *Psychological Review, 101*, 34-52.

Received July 3, 2000

Accepted July 4, 2000 ■