

UC Davis

UC Davis Previously Published Works

Title

The Roles of Motivation and Ability in Controlling the Consequences of Stereotype Suppression

Permalink

<https://escholarship.org/uc/item/5xm732fr>

Journal

Personality and Social Psychology Bulletin, 26(1)

ISSN

0146-1672

Authors

Wyer, Natalie A
Sherman, Jeffrey W
Stroessner, Steven J

Publication Date

2000

DOI

10.1177/0146167200261002

Peer reviewed

The Roles of Motivation and Ability in Controlling the Consequences of Stereotype Suppression

Natalie A. Wyer

University of California, Santa Barbara

Jeffrey W. Sherman

Northwestern University

Steven J. Stroessner

Barnard College, Columbia University

Two experiments investigated the conditions under which previously suppressed stereotypes are applied in impression formation. In Experiment 1, the extent to which a previously suppressed racial stereotype influenced subsequent impressions depended on the race of the target who was subsequently encountered. Whereas impressions of race-unspecified targets were assimilated to the stereotype following its suppression, no such effects were observed when the target belonged to the racial group whose stereotype had been initially suppressed. These results demonstrate that when perceivers are motivated to avoid stereotyping individuals, the influence of a stereotype that has been previously activated through suppression is minimized. Experiment 2 demonstrated that these processing goals effectively reduce the impact of suppression-activated stereotypes only when perceivers have sufficient capacity to enact the goals. These results suggest that both sufficient motivation and capacity are necessary to prevent heightened stereotyping following stereotype suppression.

Is stereotyping inevitable? Anyone following the growing body of research on stereotype activation might be inclined to believe that it is. Such research has demonstrated the ease with which stereotypes are activated (e.g., Banaji & Greenwald, 1995; Devine, 1989) and the wide range of consequences that follow from such activation (see Hamilton & Sherman, 1994, for a review). Indeed, for a variety of reasons, people may be inherently inclined to use their stereotypes of social groups. First, stereotypes may be efficient in that they substitute for impressions based on more effortful, individuated impression formation processes (Brewer, 1988; Fiske & Neuberg, 1990; Lippmann, 1922). Second, perceivers may boost their self-esteem by negatively stereotyping

out-group members (e.g., Fein & Spencer, 1997; Tajfel & Turner, 1986). Stereotypes also may be used to justify an unequal distribution of resources (e.g., Jost & Banaji, 1994). Finally, stereotypes may be comforting because they allow people to feel that they can predict and control their environment (e.g., Kruglanski, 1989; Lippmann, 1922).

Although a variety of factors predispose individuals to use social stereotypes, there are also numerous reasons why people may wish to limit their stereotype use. Many individuals have adopted personal standards that prohibit the use of certain stereotypes. For such individuals, using those stereotypes may lead both to feelings of guilt and to subsequent efforts to compensate for their transgressions (Monteith, 1993; Monteith, Devine, & Zuerlein, 1993). In addition, pervasive social norms discourage stereotype use in many contexts, and those who openly espouse stereotypical views may be subject to both social and legal sanctions for doing so. Awareness of norms against stereotyping may encourage even preju-

Authors' Note: This research was supported in part by a National Institute of Mental Health grant (Grant No. MH 55037) to the second author. Portions of this research were presented at meetings of the American Psychological Society (New York, June 1995) and the Midwestern Psychological Association (Chicago, May 1996). We thank Galen Bodenhausen, David Hamilton, Sarah Queller, and Jim Sherman for their helpful comments on an earlier draft. Correspondence concerning this article should be addressed to Jeffrey W. Sherman, Department of Psychology, Northwestern University, 2029 Sheridan Rd., Evanston, IL 60208-2710; e-mail: sherm@nwu.edu, or to Steven J. Stroessner, Department of Psychology, Barnard College, Columbia University, 3009 Broadway, New York, NY 10027-6598; e-mail: ss233@columbia.edu.

PSPB, Vol. 26 No. 1, January 2000 13-25

© 2000 by the Society for Personality and Social Psychology, Inc.

diced individuals to refrain from overtly expressing their stereotypes (Monteith, Deneen, & Tooman, 1996).

Although people may be motivated to suppress their stereotypes, empirical evidence suggests that attempts to do so may be unsuccessful. In fact, when people try to avoid or suppress a thought, that very thought subsequently is more likely to come to mind than if suppression was never attempted (e.g., Wegner, 1994). This phenomenon was first demonstrated by Wegner, Schneider, Carter, and White (1987) in a series of studies in which participants were instructed either to think about or to avoid (i.e., suppress) thinking about a white bear. The basic finding from these studies was that participants who initially suppressed thoughts of a white bear reported thinking about it more often in a later expression phase compared with participants who had been allowed to think about white bears all along. Thus, attempts to suppress thoughts of white bears increased the likelihood that those thoughts would later come to mind.

Wegner and his associates have proposed two mechanisms that may be used to account for the increased accessibility of suppressed thoughts. According to Wegner and Schneider's (1989) context-association model, this heightened accessibility is the result of associations formed during suppression between the unwanted (suppressed) thought and other thoughts used as distracters. After suppression has ended, these distracter thoughts may serve as retrieval cues for the formerly suppressed thought, leading to a heightened accessibility of the unwanted thought.

More recently, Wegner and Erber (1992; Wegner, 1994) have proposed an ironic process model that may also explain the heightened accessibility of suppressed thoughts. This model is based on the idea that two processes operate during thought suppression. The first process is described as an automatic monitoring process and is initiated with the intent to engage in thought suppression. This process serves to detect any occurrence of the unwanted thought so that it may be suppressed. One unintended effect of this monitoring process is that it increases the accessibility of the unwanted thought. This increase occurs for two reasons. First, the mere act of searching for a match between the to-be-suppressed thought and ongoing mental activity depends on the accessibility of that thought. If the thought were not accessible, identifying whether it had arisen would be problematic. Second, if any of the to-be-suppressed thoughts are actually detected by the monitoring process, attention is drawn to those thoughts, increasing their activation.

Once an unwanted thought is detected, a second process, described as a controlled operating process, is initiated. This process seeks out mental contents that are

more consistent with the desired state of mind (e.g., distracter thoughts) so that the unwanted mental state can be eliminated. Whereas the automatic monitoring process is unconstrained by attentional capacity, the operating process is presumed to be controlled and thus requires cognitive capacity for successful operation (Bargh, 1994). As a result, the operating process is vulnerable to interference. If the operating process is interfered with and the unwanted thought cannot be replaced, then the thought activation that results from the monitoring process may lead to a particularly heightened accessibility of the unwanted thought.

The ironic process model, then, predicts that suppressed thoughts may become more accessible during the act of thought suppression. However, this model also may be used to account for increases in accessibility after suppression has ended. That is, because the unwanted thought becomes repeatedly activated during its suppression, the activation may build to the point that it carries over to influence subsequent information processing. That is, suppression may serve to prime the unwanted thought (Macrae, Bodenhausen, Milne, & Jetten, 1994).

Consequences of Stereotype Suppression

Suppression-induced stereotype accessibility. If suppressing a thought causes it to become more accessible in memory, then suppressing a stereotype ought to produce the same effect. Macrae et al. (1994) recently demonstrated that this is, in fact, the case (for a review, see Monteith, Sherman, & Devine, 1998). In the first phase of one study (Experiment 3), participants either suppressed or did not suppress their stereotypes about a social group (skinheads). As expected, suppressors showed less use of the stereotype than nonsuppressors in this initial phase. Later, these same participants engaged in a lexical decision task in which they made judgments about stereotype-related and stereotype-unrelated words. Suppression participants showed greater response facilitation to stereotype-related words than did nonsuppressors. This finding demonstrates that the stereotype was more accessible to those participants who had suppressed the skinhead stereotype than to those who were free to express it in the initial part of the experiment.

Construct accessibility and social perception. The fact that suppressing a stereotype increases its accessibility has significant implications for social perception. Accessible constructs can influence subsequent information processing in important ways (see Higgins, 1996, for a review). For example, trait constructs made accessible through various priming manipulations have been shown to have an assimilative influence on the impressions formed of subsequently encountered individuals

(Higgins, Rholes, & Jones, 1977; Srull & Wyer, 1979). Devine (1989; see also Lepore & Brown, 1997) demonstrated similar effects on impressions of racially ambiguous targets following the subliminal priming of the African American stereotype. More recently, Macrae et al. (1994) showed that stereotypes made accessible through attempted suppression influence social perception in much the same way. Participants whose stereotypes of skinheads had been made accessible through an initial suppression task reported more stereotypic impressions of a subsequently encountered skinhead (Experiment 1) and chose to physically distance themselves from the second target to a greater extent (Experiment 2) than did participants who had not previously suppressed their stereotypes.

These results suggest that the effectiveness of stereotype suppression is, at best, short-lived. Although initial attempts to inhibit the unwanted stereotype may be successful, once the explicit goal of suppression is relaxed, the heightened accessibility of the stereotype resulting from the attempted suppression will have unintended consequences. Specifically, the heightened accessibility will influence thoughts and behavior toward subsequently encountered group members in a way opposed to the initial wishes of the suppressor. That is, subsequent group members will be stereotyped to a greater degree than if suppression had never been attempted. Obviously, these processes cast doubt on our ability to effectively control the use of our stereotypes by attempting to suppress them.

Activation Versus Application of Stereotypes

Despite the unpleasant implications of Macrae et al.'s (1994) findings, there are reasons to expect that attempts at stereotype suppression will not always fail. Although attempted suppression may unintentionally increase the accessibility of unwanted stereotypes, those stereotypes may not always be applied toward subsequently encountered group members. Several recent studies highlight the fact that activated constructs are not necessarily applied to subsequent perceptions. Sedikides (1990), for example, showed that communication goals can override accessibility effects in impression formation. In that study, participants engaged in a task that primed either positive or negative trait attributes. Subsequently, they were asked to read about an ambiguous target person and to form an impression of the person. Before reading about the target, participants were informed that they would later be asked to communicate their impressions of the target to another person (the recipient) who was described as having a positive, negative, or neutral impression of the target person. The results demonstrated standard priming effects when the

recipient was described as having a neutral attitude toward the impression target: Participants liked the target more when they were primed with positive traits than when they were primed with negative traits. The priming manipulation, however, had no effect when the recipient was described as having a positive or negative impression of the target. Regardless of the priming manipulation, participants with a positive recipient formed positive impressions of the target, and participants with a negative recipient formed negative impressions of the target. Thus, participants' communication goals obscured any effects of the activated traits.

Similar findings were reported by Thompson, Roman, Moskowitz, Chaiken, and Bargh (1994, Experiment 1; see also Ford & Kruglanski, 1995). After being primed by different trait constructs, participants were asked to form an impression of an ambiguously described target person. Participants were given either a high or low motivation to form accurate impressions. When participants had low motivation to be accurate, their impressions assimilated to the activated trait constructs. When participants were motivated to be accurate, however, their impressions were not influenced by the priming manipulation.

Thus, activated constructs clearly do not always influence later social perceptions. Personal motivations or goals can moderate the extent to which activated constructs are subsequently applied. If motivations or processing goals are incompatible with the application of activated material, then that material will not be used.

Suppression of Socially Sensitive Stereotypes: Motivational Effects

The finding that activated constructs are not always applied in later contexts may be important for our understanding of the effects of stereotype suppression. Macrae et al. (1994) demonstrated that previously suppressed stereotypes of skinheads rebounded to influence subsequent impressions of and behavior toward skinhead targets. It is important to note, however, that no strong personal or social norms existed against stereotyping skinheads as hostile for Macrae et al.'s British participants (as acknowledged by Macrae et al., 1994). Therefore, without explicit suppression instructions, participants were not likely to be motivated to avoid stereotyping the second skinhead target. As a result, the stereotype made accessible by the initial suppression task was used in forming impressions of and in guiding behavior toward the second skinhead.

A rather different set of motivations is likely involved in the suppression and subsequent application of potentially sensitive and controversial beliefs such as racial stereotypes. Many individuals have personally rejected racial stereotypes based on their belief that racial stereo-

typing is undesirable (Devine, 1989; Monteith, 1993). Furthermore, the current social climate discourages the use of racial stereotypes, a pressure that may be particularly pronounced among young, European American college students (Judd, Park, Ryan, Brauer, & Kraus, 1995). These motivations are clearly at odds with the application of accessible racial stereotypes. Therefore, based on the research of Sedikides (1990) and Thompson et al. (1994), one might expect to find very different results following the suppression of racial stereotypes than were demonstrated by Macrae et al. (1994). In particular, racial stereotypes made accessible through an initial suppression task may not be applied to subsequently encountered members of that racial group. Even with no explicit instructions to suppress stereotypes about these subsequent group members, personal and social norms may encourage participants to avoid stereotyping these individuals. Because this processing goal is at odds with the application of the previously activated stereotypes, no assimilation effects should occur (Ford & Kruglanski, 1995; Sedikides, 1990; Thompson et al., 1994). This analysis proposes an important limitation on the types of effects demonstrated by Macrae et al. (1994) and suggests that stereotypes may be controlled successfully with sufficient motivation.

EXPERIMENT 1

Overview

Experiment 1 was designed to test these hypotheses about the consequences of stereotype suppression for subsequent perception. The paradigm used in this study closely resembled the one used by Macrae et al. (1994). Participants engaged in what they believed were two separate experiments that were actually two phases of the same experiment. In the first phase, participants were presented with a picture of an African American or an Asian American male target and were asked to write a story describing a typical day in the target's life. Some participants were instructed to avoid using stereotypes as they wrote their stories. In the second phase, participants read a story about a new target whose race either was unspecified or was the same as the race of the target in the initial phase of the experiment. The target's behavior was ambiguous but could be interpreted as consistent with the stereotype of African Americans or Asian Americans. Participants formed an impression of the second target and provided ratings of him on stereotypic, counterstereotypic, and stereotype-unrelated traits.

Predictions

We expected that suppressing racial stereotypes in the initial phase of the experiment would increase the acces-

sibility of those stereotypes. Whether the accessible stereotype would be applied to the second target, however, would depend on the race of that target. In the condition in which the second target's race was unspecified, we expected that the accessible stereotype would be influential in forming an impression of the target. Because there obviously are no racial stereotypes existing for targets whose race is unspecified, there would be no basis for participants to attempt to form nonstereotypical impressions of these targets by avoiding use of the accessible material. As a result, the suppression-activated stereotypes (regarding intelligence and hostility) would be used to help interpret the target's ambiguous behaviors (which reflected varying degrees of intelligence and hostility), as in other priming experiments (e.g., Ford & Kruglanski, 1995; Sedikides, 1990; Thompson et al., 1994). Thus, we expected that impressions of the race-unspecified second targets would be more consistent with the stereotype of the initially encountered target when participants were asked to suppress the initial stereotype than when they were not. These conditions were expected to provide a conceptual replication of the findings of Macrae et al. (1994).

In contrast, in the conditions in which the second target's race matched the race of the initial target, we expected the initial suppression manipulation not to affect impressions. Although the suppression instructions would still activate the stereotypes, those stereotypes would not be applied to subsequently encountered members of the racial group. In these conditions, the target's race may serve as a cue that avoiding stereotype use is still a relevant goal. Thus, we expected that participants might be motivated to avoid stereotyping the African American and Asian American second targets (Judd et al., 1995). Because this processing goal is inconsistent with the application of the activated stereotype, the stereotype would not be used to form impressions (Ford & Kruglanski, 1995; Sedikides, 1990; Thompson et al., 1994). Consequently, impressions of the same-race second targets were expected to be equally stereotypical for suppressors and nonsuppressors.

Method

Participants. Participants were 124 undergraduates. All participants were enrolled in an introductory psychology course at the University of California, Santa Barbara (UCSB), and received partial course credit in exchange for their participation. Participants were run in groups of 2 to 6 per experimental session. No African American or Asian American participants were included in conditions in which the stereotype of their group was involved.

Design. A $2 \times 2 \times 2$ between-participants design was used in this experiment. The independent variables in

the study were the race of the target in the initial phase of the experiment (African American or Asian American replication), instructions during the initial task (suppression or no suppression), and the race of the target in the second phase of the study (same as the race of the target in the first phase or unspecified).

Procedure. The experiment consisted of two separate phases described to participants as two different experiments. To lend credibility to this cover story, each phase of the study was conducted by a different experimenter and took place in a different room.

In the initial phase of the experiment, participants were told that the study examined how people form impressions of other people at first sight. They were instructed to write a story describing a typical day in the life of a person whose photograph they would be shown. Participants in the suppression condition were informed that because past research has shown that first impressions are often influenced by cultural stereotypes about different ethnic groups, they were to avoid using such stereotypes when writing their stories. Participants in the nonsuppression condition were given no such special instructions. All participants were then given 5 minutes to compose their stories about the person.

After writing their stories, participants were taken to a different room in which a second experimenter gave instructions for the second phase of the experiment. They were informed that the second experiment was being conducted to study how people form impressions of others based on relatively little information. Participants read a story about a target named Robert and were asked to form an impression of Robert as they read the story. In the story, Robert was described as engaging in three hostile behaviors, three unintelligent behaviors, three passive behaviors, and three intelligent behaviors (see appendix). Thus, the story contained equal amounts of information that was stereotypic of African Americans (i.e., the hostile and unintelligent behaviors; see Devine, 1989; Devine & Baker, 1991, for traits stereotypic of African Americans) and stereotypic of Asian Americans (i.e., the passive and intelligent behaviors; see Jackson, Lewandowski, Ingram, & Hodge, 1997; Rothbart & John, 1993, for traits stereotypic of Asian Americans), causing Robert to be ambiguous along these dimensions. The content of the stories was identical across conditions, except for Robert's race. In the unspecified race condition, no mention was made of Robert's race. In the same-race condition, Robert's race was identified as the same as that of the person about whom participants had written in the first phase of the experiment (e.g., "Robert is a 25-year-old African American/Asian American male from Southern California").

After reading the story about Robert, participants were asked to rate him along 12 personality trait dimen-

sions. These trait dimensions included several stereotype-relevant traits (hostile, aggressive, passive, respectful, smart, and unintelligent; see Devine, 1989; Devine & Baker, 1991; Jackson et al., 1997; Rothbart & John, 1993) and several stereotype-irrelevant traits (foolish, friendly, unpleasant, intellectual, likable, boring). Participants rated Robert on each of these dimensions using 9-point Likert-type scales (1 = *not at all*, 9 = *extremely*).

Results

Manipulation check. To ensure that participants in the suppression condition actually suppressed their stereotypes, the stories written during the initial phase of the experiment were rated for stereotypic content. Two coders who were blind to participants' condition (including the race of the target) rated each story on a 9-point scale regarding its consistency with both the African American and the Asian American stereotype. The coders' ratings were quite consistent, $r = .796$, $p < .01$. An analysis of variance (ANOVA) of the average of their stereotypicality ratings yielded a significant main effect of suppression instructions, $F(1, 123) = 4.75$, $p < .05$. As predicted, participants in the suppression condition wrote stories that were significantly less stereotypical of the target's racial group ($M = 4.15$) than did participants in the nonsuppression condition ($M = 5.03$). Thus, participants in the suppression condition suppressed their stereotypes during Phase 1 of the experiment.

Stereotypicality of trait ratings. To analyze the stereotypicality of trait ratings made about the target in the second phase of the experiment, a stereotype composite measure was computed. For those participants who read about the Asian American target in Phase 1, their ratings of the second target on the traits "hostile" (reverse scored), "unintelligent" (reverse scored), "aggressive" (reverse scored), "passive," "respectful," and "smart" were averaged. For participants who read about the African American target, their ratings on the same traits were averaged, except that only the latter three ratings were reverse scored. Thus, higher ratings for both targets reflected greater stereotypicality.

These composite measures were subjected to a three-way (Initial Target Race Replication \times Initial Task Instructions \times Race of the Second Target) ANOVA. We had predicted that there would be a significant two-way interaction between the initial task instructions and whether the race of the second target was specified. Specifically, among participants for whom the race of the second target was unspecified, we expected to see significantly greater stereotype use in the suppression condition than in the nonsuppression condition. In contrast, we expected there to be no differences in stereotype use among participants for whom the race of the second target matched the race of the initial target. We did not

expect the target race replication to qualify the predicted effect, and it did not, $F(1, 116) < 1$ for the three-way interaction. Although the pattern of cell means was consistent with the predicted two-way interaction (see Figure 1), the ANOVA results did not reach standard levels of statistical significance, $F(1, 116) = 1.59, p > .2$.

However, our hypotheses were clear regarding the pattern of participants' responses depending on whether the race of the second target was specified. Specifically, in the condition in which race was not specified, we predicted priming effects in the suppression condition, which would be reflected by more stereotypical ratings of the target. An analysis comparing the ratings in the race unspecified condition was significant, $F(1, 60) = 5.77, p < .05$. Consistent with our predictions, participants who had initially suppressed the target's stereotype rated this target more consistently with the stereotype ($M = 5.17$) than did participants who had not suppressed the stereotype ($M = 4.65$). A different pattern of results was predicted when participants were informed that the second target's race matched that of the initial target. Specifically, we predicted that priming effects induced by stereotype suppression would not occur in this case. Consistent with our expectations, whether participants had suppressed the stereotype did not influence their ratings of the second same-race target. In particular, participants who suppressed ($M = 5.07$) and did not suppress ($M = 4.99$) the stereotype did not differ in the stereotypicality of their ratings, $F(1, 56) < 1$ (see Figure 1).

Thus, participants in the unspecified race condition who suppressed the stereotype of either African Americans or Asian Americans were more likely to evaluate the second phase target's behavior as consistent with the suppressed stereotype than participants who did not suppress the stereotype. Participants in the conditions in which the second target was the same race as the first target showed a different pattern; they showed the same level of stereotyping regardless of suppression instructions. These data are consistent with the notion that whereas suppression increases the accessibility of stereotypes, the consequences of that suppression depend on whether the subsequent targets' race is specified.

Readers may note that ratings in the suppress/race unspecified condition ($M = 5.17$) were not more stereotypical than participants in the suppress/race specified ($M = 5.07$) and no suppress/race specified ($M = 4.99$) conditions. This might seem to argue against our conclusion that rebound effects were occurring in the race-unspecified conditions. However, there are good reasons to reject this logic. First, the targets in the race specified and unspecified conditions are fundamentally different targets. Whereas no race information was provided in the unspecified conditions, the targets in the race specified conditions were described as either African Ameri-

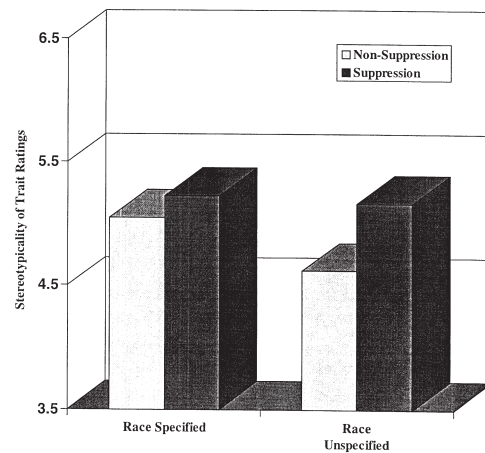


Figure 1 Ratings of Phase 2 targets on stereotypical traits as a function of suppression instructions and race specification, Experiment 1.

can or Asian American. This diminishes the utility of comparisons between targets because it is clearly evident from a comparison of the no suppress conditions. In those conditions, the race specified (Asian American and African American) targets were rated as much more stereotypical (4.99) than the race unspecified target (4.65). Given that racial information was provided in the former but not the latter case, this is not terribly surprising. Thus, what this means is that, at baseline, the specified targets were already seen as more stereotypical than the unspecified targets. However, what is most important is that ratings of the specified targets did not become even more stereotypical following suppression (4.99 vs. 5.07). In contrast, ratings of the unspecified target were much more stereotypical following suppression (4.65 vs. 5.17). This increase in stereotyping of the unspecified target reflects a clear rebound effect. Nonetheless, even though the ratings of the unspecified target were more stereotypical in the suppression than the nonsuppression condition, they were not more stereotypical than ratings of the specified targets because the specified targets were seen as more stereotypical to begin with. This highlights the more general point that investigations of thought suppression that rely on target ratings for evidence of rebound effects should be careful to focus on comparisons of the same target following suppression and nonsuppression.

Readers also may note that we did not collect baseline ratings of the secondary target from participants who did not engage in the first phase of the experiment (i.e., there were no participants who did not write "day in the

life" stories). This baseline would have permitted us to examine the extent of stereotype activation among suppressors as compared to people who had no initial encounter with a member of a stereotyped group. However, the goal of this research was to investigate stereotype activation following an initial encounter with a member of a stereotyped group. In particular, if a person encounters a member of a stereotyped group and suppresses the stereotype, will the stereotype be more accessible than if he or she encounters a member of the group and does not suppress the stereotype? As such, the baseline we collected (i.e., see a member of the group but do not suppress the stereotype) was the most appropriate baseline. Moreover, the "no suppress" baseline we collected is more conservative than a "no encounter" baseline. Although most studies of stereotype suppression have not included a "no encounter" baseline (for a review, see Monteith et al., 1998), Macrae et al. (1994, Experiment 3) showed that simply encountering a member of a stereotyped group activated the stereotype in comparison to a "no encounter" baseline. Thus, the fact that suppressing the stereotype produced even greater accessibility than baseline encounter activation in the present study is telling, indeed.

Discussion

The results from the race-unspecified second target conditions provide further evidence that attempts to suppress stereotypes increase the accessibility of those stereotypes. Participants in these conditions who suppressed the stereotype of either African Americans or Asian Americans were more likely to evaluate the second phase target's behavior as consistent with the suppressed stereotype than participants who did not suppress the stereotype. These results are analogous to those reported by Macrae et al. (1994). However, the results from the same-race second target conditions demonstrate that these suppression-activated stereotypes will not necessarily be applied to impressions of subsequently encountered group members. Participants in these conditions showed the same level of second target stereotyping, regardless of suppression instructions. Thus, the unintended postsuppression stereotyping effects demonstrated by Macrae et al. (1994) may not occur when perceivers have inhibited socially sensitive stereotypes.

One key factor that seems to determine whether suppression-activated stereotypes will be applied to other group members relates to the processing goals of the perceiver when learning about those group members. The use of skinhead targets in Macrae et al.'s (1994) research ensured that participants would not be concerned about forming stereotypical impressions of the postsuppression targets because it is normative to

use those stereotypes. As a result, the suppression-activated stereotype was applied to judgments of the second skinhead target. Similarly, in our race-unspecified conditions, participants had no reason to attempt to form nonstereotypic impressions of the second target. Because no race information was given for the second target, using the activated stereotypes would not be perceived to result in stereotyping of the second target. Thus, avoiding stereotyping was not a concern for participants in the race-unspecified conditions, and therefore, the previously activated stereotypes influenced these impressions.

In contrast, given the prevailing personal and social norms against racial stereotyping, participants in our same-race second target conditions were likely motivated to avoid stereotyping the second targets. This goal may have been particularly salient for suppression participants given that antistereotyping norms had already been made salient by the initial suppression instructions. As a result, the suppression-activated stereotypes did not influence the impressions formed of subsequent same-race targets. This finding is consistent with previous demonstrations that activated constructs will not be applied in social perception when their application is inconsistent with the current processing goal of the perceiver (Ford & Kruglanski, 1995; Sedikides, 1990; Thompson et al., 1994). This suggests that perceivers' motivations play an important role in determining the effectiveness of stereotype suppression.

EXPERIMENT 2

Motivation, however, is not the only factor that determines whether an activated construct will be subsequently applied. For a motivated processing goal to override the effects of an activated construct, the perceiver must have sufficient cognitive resources to achieve that processing goal. Several studies, in fact, suggest that the depletion of processing capacity increases the likelihood that activated constructs will influence subsequent information processing, even if perceivers have opposing processing goals. Thompson et al. (1994, Experiment 2; see also Ford & Kruglanski, 1995), for example, primed participants with trait constructs before giving them an impression formation task. In addition to manipulating impression accuracy motivation in this study, as was described earlier, the researchers also placed some participants under a cognitive load as they reported their impressions of the target. The results demonstrated that priming effects occurred if participants had either low accuracy motivation or limited processing capacity. Only when participants were both motivated to be accurate and had full processing capacity were impressions unaffected by the primed material. Gilbert and Hixon (1991) demonstrated a similar effect in the activation and appli-

cation of stereotypes. Participants whose stereotypes had been activated in an initial phase of an experiment were more likely to apply those stereotypes only when their capacity was depleted during the application phase. When capacity was high during the application phase, the activated stereotypes did not influence target impressions.

Thus, processing goals may diminish the influence of activated mental constructs only if significant cognitive resources are available to the perceiver. When resources are low, activated constructs may be applied to subsequent targets, regardless of perceivers' motivations. Whether manipulations of cognitive capacity would similarly influence the application of suppression-activated stereotypes was the focus of Experiment 2.

Overview and Predictions

Experiment 2 was designed to test the hypothesis that perceivers' ability to overcome the influence of suppression-activated stereotypes is subject to capacity constraints. The same-race conditions of Experiment 1 were replicated with the addition of conditions in which participants were placed under a cognitive load during the second target impression formation task. In the first phase of the experiment, participants either suppressed or did not suppress their stereotypes of African Americans or Asian Americans. Then, in a supposedly independent experiment, participants formed an impression of a second member of the same racial group. Whereas half of the participants were required to rehearse an eight-digit number while engaging in the second impression formation task, the others were not given this additional task. This task has been shown to effectively diminish processing resources (e.g., Gilbert & Hixon, 1991; Thompson et al., 1994).

We expected to replicate the results of the same-race conditions of Experiment 1 in the low cognitive load conditions of Experiment 2. The stereotypes activated by the suppression task would not be applied to impressions of the second target. Thus, suppressors and nonsuppressors would demonstrate the same amount of second target stereotyping. Different results were expected in the high load conditions. The high cognitive load should interfere with suppressors' ability to enact their goal of not stereotyping the second target. As a result, these participants would be unable to avoid using their activated stereotypes to interpret the behavior of the second target, and their impressions would reflect the influence of the stereotype made accessible through attempted suppression (Ford & Kruglanski, 1995; Gilbert & Hixon, 1991; Thompson et al., 1994). Thus, in the high load condition, we expected that suppressors would form more stereotypical impressions of the sec-

ond targets than would nonsuppressors. These results would demonstrate that just as reducing participants' motivation to avoid stereotyping results in suppression-induced stereotyping effects (as in the race-unspecified conditions of Experiment 1) so too does interfering with their capacity to control their use of accessible stereotypes.

Method

Participants. A total of 162 UCSB undergraduates (110 females) participated in the experiment in exchange for partial course credit. Participants were run in groups of 2 to 7. No Asian American or African American students participated in conditions in which the stereotype of their own racial group was involved.

Design. The experiment entailed a $2 \times 2 \times 2$ between-participants design. Target race replication (Asian American or African American), instructions (suppression or no suppression), and cognitive load (high or low) were manipulated. The primary dependent measure consisted of participants' ratings of the second target on traits related to hostility and intelligence.

Procedure. Participants engaged in what they were told were two independent experiments. The procedures for the first phase of the experiment were identical to those used in the first phase of Experiment 1. Specifically, participants wrote a short story about a typical day in the life of a man named John (who was either an African American or Asian American) whose photograph they were shown. Whereas half of the participants were instructed to suppress their stereotypes, the others were given no special instructions.

After completing the first phase of the experiment, participants were introduced to the second experiment. They were informed that they would read a story about a man named Robert and that they should try to form an impression of him. Participants in the high load condition were further told that one purpose of the experiment was to determine how people form impressions when they engage in simultaneous tasks. To study this, participants were told that they would be required to remember an eight-digit number while they read the story. They also were told that they would be asked to recall the number at the end of the experiment. The experimenter then read the number aloud once (at the rate of approximately one digit per second). After reciting the number once, the experimenter instructed the participants to begin reading the story. Participants in the nondistraction condition were simply told to form an impression of Robert as they read the story. Always, Robert was described as being of the same race as the target in the first phase of the experiment. After reading the story, participants reported their eight-digit number

and then completed a questionnaire that included the trait scales on which they were to rate Robert. Finally, participants completed a questionnaire in which they were asked about their reactions to the experiment. Included in the questionnaire were items that probed for suspicion regarding the independence of the two phases of the experiment.

Results

Manipulation checks. As in Experiment 1, the stories written during the initial phase of the experiment were rated for stereotypic content. Two coders who were blind to participants' condition (including the race of the target) rated each story on a 9-point scale as to its consistency with both the African American and the Asian American stereotype. Their ratings were sufficiently consistent ($r = .701, p < .01$); thus, the average of their ratings was used in the subsequent analysis. An ANOVA on the stereotypicality ratings yielded a significant main effect of suppression instructions, $F(1, 162) = 8.59, p < .01$. As predicted, participants in the suppression condition wrote stories that were significantly less stereotypic of the target's racial group ($M = 5.52$) than did participants in the nonsuppression condition ($M = 5.97$). Thus, participants in the suppression condition suppressed their stereotypes during Phase 1 of the experiment.

Recall of the eight-digit number was examined as a manipulation check on the cognitive capacity manipulation. Participants who failed to recall at least half of the digits of the eight-digit number correctly were deemed to have not rehearsed the number diligently enough to reduce available cognitive capacity (see Gilbert & Hixon, 1991; Thompson et al., 1994). On this basis, the data from 17 participants were excluded from the analyses.

Finally, we examined participants' responses to the questionnaire that assessed their level of suspicion regarding the independent nature of the two experimental tasks. Although a number of participants mentioned that both experiments were concerned with impression formation and that they both involved minority targets, none of the participants expressed suspicion that the two phases were actually part of the same experiment.² Thus, the explicit suppression demands of the first phase of the experiment were not perceived to apply during the second phase.

Stereotypicality of trait ratings. As in Experiment 1, the stereotypicality of trait ratings of the second target was assessed based on a stereotype composite measure. Higher scores on the composite measure always reflected greater stereotypicality.

These composite measures were subjected to a 2 (target race replication) \times 2 (suppression) \times 2 (cognitive load) ANOVA. This analysis yielded a main effect of target race, $F(1, 137) = 95.37, p < .001$, but target race did

not qualify the predicted two-way interaction, $F(1, 137) < 1$ for the three-way interaction. The predicted two-way interaction between suppression and cognitive load was significant, $F(1, 137) = 4.15, p < .05$. This interaction is illustrated in Figure 2. There was no effect of suppression instructions for participants in the low cognitive load condition, $F(1, 75) < 1$. Participants who suppressed their stereotypes produced ratings that were no more stereotypic ($M = 4.88$) than those produced by participants who did not suppress ($M = 5.04$). This finding replicates the pattern of results obtained in the same-race condition of Experiment 1. However, there was a significant effect of suppression instructions for participants in the high cognitive load condition, $F(1, 62) = 4.13, p < .05$, such that participants in the suppression condition rated the second phase target more stereotypically ($M = 5.40$) than did those in the nonsuppression condition ($M = 4.73$). Thus, suppressing racial stereotypes affected the subsequent evaluations made by participants in the high cognitive load condition but did not affect judgments of participants in the low cognitive load condition.

Discussion

The low cognitive load conditions of Experiment 2 replicated the same-race conditions of Experiment 1. Suppression-activated stereotypes did not influence the impressions formed of subsequently encountered members of the targeted group. Thus, suppressors and nonsuppressors formed equally stereotypical impressions of the second phase targets. In contrast, suppressors in the high cognitive load condition were unable to override the influence of their activated stereotypes and formed more stereotypical impressions of the second target than did nonsuppressors. These results demonstrate that motivation alone is insufficient for avoiding the application of suppression-activated stereotypes to other group members. The motive to avoid stereotyping those group members can override such effects but only when sufficient capacity is available to enact that motive.

GENERAL DISCUSSION

The results of these experiments demonstrate when and why stereotype suppression does and does not lead to increased stereotyping. In Experiment 1, participants applied activated stereotypes when the targets whom they encountered in the second phase were race-unspecified. In this case, participants were not motivated to form nonstereotypical impressions of the second targets because there is no reason to fear stereotyping a target whose race is not known. In Experiment 2, suppressed stereotypes were applied in the impression formation process when cognitive resources were depleted, even though the second target belonged to the same racial category as the initial target. In combination,




Figure 2 Ratings of Phase 2 targets on stereotypic traits as a function of suppression instructions and cognitive load, Experiment 2.

these results demonstrate that a disruption either in perceivers' motivation to avoid stereotyping or in their capacity to control the use of an activated stereotype will lead to increased use of a stereotype following its suppression. When both motivation and capacity to avoid applying a suppression-activated stereotype are present, however, those stereotypes will not influence subsequent impressions. This result was demonstrated in the same-race conditions of Experiment 1 and was replicated in the low cognitive load conditions of Experiment 2. Thus, stereotype suppression does not lead inevitably to greater subsequent use of that stereotype. These results extend the findings of recent research investigating the conditions under which activated constructs are and are not applied (Ford & Kruglanski, 1995; Gilbert & Hixon, 1991; Newman, Duff, Hedberg, & Blitstein, 1996; Sedikides, 1990; Thompson et al., 1994).

The results of our same-race conditions may be contrasted with those reported by Macrae et al. (1994). Whereas increased stereotyping of subsequently considered group members following suppression of the group stereotype was observed in those studies, such effects did not emerge in the same-race conditions of the present experiments (save for the high load condition in Experiment 2). The primary difference between the same-race conditions in our experiments and Macrae et al.'s experiments was that targets in our experiments belonged to social groups about which there are strong social, and often personal, standards against stereotyping. Participants in Macrae et al.'s experiment were

unlikely to be concerned with negatively stereotyping skinheads when they encountered a second target from the same group because there are no strong norms against stereotyping skinheads. In fact, negatively stereotyping such groups may be quite normative. Consequently, participants' impressions of the second target were open to the influence of a suppression-activated stereotype. In contrast, personal beliefs against racial stereotyping are widely held (this belief is almost uniform among European American college students) (Judd et al., 1995), and there are strong social norms against doing so. Thus, in contrast to Macrae et al.'s participants, it is likely that participants in our experiments were motivated to avoid stereotyping the second phase targets. This motive is incompatible with the application of group stereotypes regardless of whether such stereotypes have been made more accessible through prior suppression. Consistent with other research (Ford & Kruglanski, 1995; Sedikides, 1990; Thompson et al., 1994), the presence of this incompatible goal obscured any potential impact of the activated stereotype.

Implications for the Inevitability of Stereotyping

The results from the same-race conditions of our experiments suggest that if perceivers have sufficient motivation and sufficient cognitive resources they can avoid using accessible stereotypes to judge other group members whom they encounter. On the surface, this finding seems to paint a rather optimistic picture of our ability to control our stereotypes. One issue for concern, however, regards whether individuals will generally have sufficient cognitive resources available to engage in the processes that appear necessary to avoid stereotype use. The results of Experiment 2 suggest that the presence of even a minor distraction is enough to disrupt individuals' ability to avoid the use of suppression-activated stereotypes in forming their impressions. Given that everyday life is full of distractions, the outlook for individuals successfully controlling their use of activated racial stereotypes does not seem so bright. In addition to simple distractions, affective states such as positive moods (Bodenhausen, Kramer, & Susser, 1994; Mackie, Queller, Stroessner, & Hamilton, 1996; Stroessner, Hamilton, & Mackie, 1992; Stroessner & Mackie, 1992) and anxiety (Wilder & Shapiro, 1989) have been shown to decrease the ability to engage in deliberative processing. Likewise, physical fatigue also decreases available processing capacity (e.g., Bodenhausen, 1990; Kim & Baron, 1988). Therefore, given that an individual's cognitive capacity varies tremendously in any day, stereotyping following suppression may be more likely to occur in naturalistic settings than in the lab.

Questions of Motivation

Although we did not directly manipulate motivation in this research, there is ample reason to believe that participants were motivated to avoid using their stereotypes in the second phases of both Experiment 1 (in the same-race conditions) and Experiment 2. A considerable amount of research has demonstrated that people attempt to curtail their use of stereotypes in situations, such as those in the current experiment, that make social norms against stereotyping salient (e.g., Fazio, Jackson, Dunton, & Williams, 1995; Wyer, Sherman, & Stroessner, 1998). Furthermore, although there are undoubtedly individual differences in the extent to which individuals are personally motivated to avoid stereotyping (see below), research by Judd et al. (1995) suggests that young, European American college students are particularly likely to be so motivated. They also may be more sensitive to salient social norms to avoid stereotyping. These prior findings and the results of this experiment lead us to conclude that motivation played a crucial role in allowing participants in the same-race conditions to overcome the stereotype accessibility that occurred as a result of suppression.

There are a number of ways that the motive to avoid stereotyping the same-race second targets could have undermined stereotype application. One possibility is that suppression participants reinitiated suppression attempts when they encountered the second group member, masking any lingering accessibility of the stereotype resulting from the initial suppression attempt. However, this account cannot easily explain the data. If the suppressors had intentionally suppressed their stereotypes about the second target, then their judgments of that target would have been less stereotypical than those of the nonsuppressors, as was the case with the impressions that were formed in the initial phase of the experiment. They were not. A second possibility is that the effects of the stereotype being primed during the first phase of the experiment and the effects of suppression being reactivated in the second phase of the experiment acted in opposition to each other. That is, both processes were active in the suppressors but they canceled each other out, resulting in no differences between the suppression and nonsuppression conditions. Alternatively, it may be that suppressors varied in the extent to which suppression demands were reactivated. That is, some participants may have reinstigated the suppression process and others did not. These two types of participants, showing opposite effects, would have produced the observed pattern of data. Arguing against this possibility, however, is the fact that the variance in judgments in the suppression and no-suppression conditions did not differ. Finally, and we believe most likely, participants in the same-race conditions may have been moti-

vated by egalitarian social and personal norms to form individuated impressions of the second target that were based on the specific behaviors described in the story. Because the information provided about the target was mixed with regard to its consistency with the suppressed stereotype, individuated evaluations of the target would be relatively independent of the stereotype for both suppressors and nonsuppressors.

This analysis suggests that sufficient motivation to avoid stereotyping is necessary for individuals to avoid the ironic effects of stereotype suppression. One important implication of this point concerns the fact that individuals differ in their level of chronic motivation to avoid stereotyping (Dunton & Fazio, 1997; Monteith et al., 1993; Plant & Devine, 1998). For example, to the extent that certain individuals hold personal standards against expressing prejudice, they may be chronically motivated to avoid stereotyping others. By contrast, more highly prejudiced individuals may only attempt to avoid stereotyping if the social situation demands it (Devine, Monteith, Zuwerink, & Elliot, 1991; Dunton & Fazio, 1997; Monteith et al., 1993; Plant & Devine, 1998). As a result, these persons may be unlikely to ever avoid the influence of a suppression-activated stereotype.

These different motivations to avoid stereotyping also may have important implications for people's ability to overcome the influence of previously activated (through suppression or otherwise) stereotypes. In particular, through practice at activating egalitarian beliefs and forming individuated impressions, those who are internally motivated to avoid stereotyping may have become quite adept at this task. As a result, these individuals may be in a relatively strong position to avoid the influence of a previously activated stereotype in subsequent encounters with group members, even if capacity is limited. Thus, they may simply be efficient at stereotype avoidance (see Monteith et al., 1998, for a review). In contrast, those who are motivated by external forces to avoid stereotyping may have relatively little experience doing so. As a result, previously activated stereotypes may be more likely to impinge on subsequent encounters with group members, even in the rare case that the perceiver is trying to avoid stereotyping. This may be particularly true if the perceiver's resources are somehow depleted. Clearly, one important goal for future research will be to examine how individual differences in the motivation to avoid stereotyping may moderate the repercussions of stereotype suppression.

Conclusion

The present research provides compelling evidence that the relation between stereotype suppression and stereotype use is complex. Both motivation and ability to avoid stereotyping appear necessary for successful

stereotype suppression to be achieved without the risk of unintentional consequences. Given that stereotype suppression is an activity in which people frequently engage (Bodenhausen & Macrae, 1996; Fazio et al., 1995; Monteith, 1993; Monteith et al., 1998; Wyer et al., 1998), it will be important to identify the conditions that spontaneously generate attempts at suppression and additional factors that will determine whether suppression will be successful. That suppression increased subsequent stereotyping in some conditions of our experiments does not suggest that stereotype suppression is necessarily bad. On the contrary, in many situations, inhibiting stereotypic thinking is critical. We should be aware, however, that these efforts may influence our subsequent social perceptions and behaviors in important and unexpected ways and may occasionally backfire if we lose the motivation or the ability to correct for a suppression-activated stereotype.

APPENDIX

Robert is a 25-year-old (African American/Asian American) from Southern California. He lives in an apartment with his roommate Jim.

Robert enjoys doing various things in his free time. He reads a lot and likes to listen to music. In fact, Robert has been having some trouble with some of the other people in his apartment building because he plays his stereo much too loud and refuses to turn it down when his neighbors complain.

Robert recently went into the city to go to the art museum because there was an exhibit that he wanted to see. He had a little trouble getting there, however, because he got lost and was confused by the freeway map, so he had to stop at a gas station to ask for directions. As he left the station, Robert heard the attendant insult him under his breath, but Robert did not turn around to confront him.

On his way back from the museum, Robert's car broke down. After calling for a tow truck, Robert found the nearest bus stop so that he could catch the next bus home. Because it was late in the afternoon, there were a lot of people who also wanted to catch a bus out of the city. When the bus finally arrived, Robert shoved his way to the front so that he could be sure that he would be able to get on.

The next day at work, Robert had a meeting in the morning. Everyone at the company was meeting so that they could brainstorm some new ideas for improving performance during the upcoming year. Although Robert had a good idea, he did not mention it because he did not want to interrupt the other people that were talking.

Later that day, Robert was working on the computer when he accidentally erased all of the word processing files. Luckily, his supervisor kept a back-up copy of the files.

When Robert got home from work, his roommate's father had just arrived to visit for a couple of days. The three of them went out to dinner at a local restaurant where Robert got into

a long conversation with his roommate's father about foreign affairs.

After dinner, Robert left his roommate and his father to stop by the drug store to pick something up. The store was very busy and Robert had to wait in line for almost a half hour. Although he was a little upset by this, he did not object when the person in front of him allowed three friends to get in line in front of him.

When Robert got home, he found that his roommate had left some dirty dishes in the kitchen sink. Annoyed by this, Robert wrote a nasty note to his roommate reminding him to clean up after himself better.

Robert decided that he would watch some television before going to bed. He had trouble getting the television to work and had to fool around with it for several minutes before he realized that the set was unplugged. Robert then watched an old sitcom for a few minutes before falling asleep.

NOTES

1. A breakdown of participants' sex is not available for Experiment 1. Participant sex did not have any effects in Experiment 2.

2. This type of questionnaire was not included in Experiment 1. However, Experiment 2 involved the conditions of Experiment 1 in which participants could be expected to be most suspicious about the independence of the two phases (i.e., in which the target individuals in both phases were the same race). Thus, we can infer that participants in Experiment 1 were likely to be no more suspicious than those in Experiment 2.

REFERENCES

- Banaji, M. R., & Greenwald, A. G. (1995). Implicit gender stereotyping in judgments of fame. *Journal of Personality and Social Psychology*, *68*, 181-198.
- Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition, Vol. 1: Basic processes* (pp. 1-40). Hillsdale, NJ: Lawrence Erlbaum.
- Bodenhausen, G. V. (1990). Stereotypes as judgmental heuristics: Evidence of circadian variations in discrimination. *Psychological Science*, *1*, 319-322.
- Bodenhausen, G. V., Kramer, G., & Susser, K. (1994). Happiness and stereotypical thinking in social judgment. *Journal of Personality and Social Psychology*, *66*, 621-632.
- Bodenhausen, G. V., & Macrae, C. N. (1996). The self-regulation of intergroup perception: Mechanisms and consequences of stereotype suppression. In C. N. Macrae, C. Stangor, & M. Hewstone (Eds.), *Stereotypes and stereotyping* (pp. 227-254). New York: Guilford.
- Brewer, M. B. (1988). A dual process model of impression formation. In T. K. Srull & R. S. Wyer, Jr. (Eds.), *Advances in social cognition* (Vol. 1, pp. 1-36). Hillsdale, NJ: Lawrence Erlbaum.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 5-18.
- Devine, P. G., & Baker, S. M. (1991). Measurement of racial stereotype subtyping. *Personality and Social Psychology Bulletin*, *17*, 44-50.
- 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.
- 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.
- Fein, S., & Spencer, S. J. (1997). Prejudice as self-image maintenance: Affirming the self through derogating others. *Journal of Personality and Social Psychology*, 73, 31-44.
- Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. In M. P. Zanna (Ed.), *Advances in Experimental Social Psychology* (Vol. 23, pp. 1-74). New York: Academic Press.
- Ford, T. E., & Kruglanski, A. W. (1995). Effects of epistemic motivations on the use of accessible constructs in social judgment. *Personality and Social Psychology Bulletin*, 21, 950-962.
- Gilbert, D. T., & Hixon, J. G. (1991). The trouble of thinking: Activation and application of stereotypic beliefs. *Journal of Personality and Social Psychology*, 60, 509-517.
- Hamilton, D. L., & Sherman, J. W. (1994). Stereotypes. In R. S. Wyer, Jr., & T. K. Srull (Eds.), *Handbook of social cognition*, Vol. 2: Applications (pp. 1-68). Hillsdale, NJ: Lawrence Erlbaum.
- Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 133-168). New York: Guilford.
- Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology*, 13, 141-154.
- Jackson, L. A., Lewandowski, D. A., Ingram, J. M., & Hodge, C. N. (1997). Group stereotypes: Content, gender specificity, and affect associated with typical group members. *Journal of Social Behavior and Personality*, 12, 381-396.
- 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.
- Judd, C. M., Park, B., Ryan, C. S., Brauer, M., & Kraus, S. (1995). Stereotypes and ethnocentrism: Diverging interethnic perceptions of African American and White American youth. *Journal of Personality and Social Psychology*, 69, 460-481.
- Kim, H.-S., & Baron, R. S. (1988). Exercise and the illusory correlation: Does arousal heighten stereotypic processing? *Journal of Experimental Social Psychology*, 24, 366-380.
- Kruglanski, A. W. (1989). *Lay epistemics and human knowledge*. New York: Plenum.
- Lepore, L., & Brown, R. (1997). Category and stereotype activation: Is prejudice inevitable? *Journal of Personality and Social Psychology*, 72, 275-287.
- Lippmann, W. (1922). *Public opinion*. New York: Harcourt Brace Jovanovitch.
- Mackie, D. M., Queller, S., Stroessner, S. J., & Hamilton, D. L. (1996). Making stereotypes better or worse: Multiple roles for positive affect in group impressions. In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition*, Volume 3 (pp. 371-396). New York: Guilford.
- 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.
- Monteith, M. J. (1993). Self-regulation of prejudiced responses: Implications for progress in prejudice-reduction efforts. *Journal of Personality and Social Psychology*, 65, 469-485.
- Monteith, M. J., Deneen, N. E., & Tooman, G. D. (1996). The effect of social norm activation on the expression of opinions concerning gay men and Blacks. *Basic and Applied Social Psychology*, 18, 267-288.
- Monteith, M. J., Devine, P. G., & Zuwerink, J. R. (1993). Self-directed versus other-directed affect as a consequence of prejudice-related discrepancies. *Journal of Personality and Social Psychology*, 64, 198-210.
- Monteith, M. J., Sherman, J. W., & Devine, P. G. (1998). Suppression as a stereotype control strategy. *Personality and Social Psychology Review*, 2, 63-82.
- Newman, L. S., Duff, K. J., Hedberg, D. A., & Blitstein, J. (1996). Rebound effects in impression formation: Assimilation and contrast effects following thought suppression. *Journal of Experimental Social Psychology*, 32, 460-483.
- Plant, E. A., & Devine, P. G. (1998). Internal and external sources of motivation to control prejudice. *Journal of Personality and Social Psychology*, 75, 811-832.
- Rothbart, M., & John, O. P. (1993). Intergroup relations and stereotype change: A social-cognitive analysis and some longitudinal findings. In P. M. Sniderman, P. E. Tetlock, & E. G. Carmines (Eds.), *Prejudice, politics, and the American dilemma* (pp. 32-59). Stanford, CA: Stanford University Press.
- Sedikides, C. (1990). Effects of fortuitously activated constructs versus activated communication goals on person impressions. *Journal of Personality and Social Psychology*, 58, 397-408.
- Srull, T. K., & Wyer, R. S. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. *Journal of Personality and Social Psychology*, 37, 1660-1672.
- Stroessner, S. J., Hamilton, D. L., & Mackie, D. M. (1992). Affect and stereotyping: The effect of induced mood on distinctiveness-based illusory correlations. *Journal of Personality and Social Psychology*, 62, 564-576.
- Stroessner, S. J., & Mackie, D. M. (1992). The impact of induced affect on the perception of variability in social groups. *Personality and Social Psychology Bulletin*, 18, 546-554.
- Tajfel, H., & Turner, J. C. (1986). The social identity theory of intergroup behavior. In S. Worchel & W. Austin (Eds.), *The psychology of intergroup relations* (2nd ed., pp. 7-24). Chicago: Nelson-Hall.
- Thompson, E. P., Roman, R. J., Moskowitz, G. B., Chaiken, S., & Bargh, J. A. (1994). Accuracy motivation attenuates covert priming: The systematic reprocessing of social information. *Journal of Personality and Social Psychology*, 66, 474-489.
- Wegner, D. M. (1994). Ironic processes of mental control. *Psychological Review*, 101, 34-52.
- Wegner, D. M., & Erber, R. (1992). The hyperaccessibility of suppressed thoughts. *Journal of Personality and Social Psychology*, 63, 903-912.
- Wegner, D. M., & Schneider, D. J. (1989). Mental control: The war of the ghosts in the machine. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 287-305). New York: Guilford.
- Wegner, D. M., Schneider, D. J., Carter, S. R., & White, T. L. (1987). Paradoxical effects of thought suppression. *Journal of Personality and Social Psychology*, 53, 5-13.
- Wilder, D. A., & Shapiro, P. (1989). Effects of anxiety on impression formation in a group context: An anxiety-assimilation hypothesis. *Journal of Experimental Social Psychology*, 25, 481-499.
- Wyer, N. A., Sherman, J. W., & Stroessner, S. J. (1998). The spontaneous suppression of racial stereotypes. *Social Cognition*, 16, 340-352.

Received June 18, 1997

Revision accepted October 1, 1998