

UC Davis

UC Davis Previously Published Works

Title

Prejudice and Stereotype Maintenance Processes: Attention, Attribution, and Individuation

Permalink

<https://escholarship.org/uc/item/3zq8s5cq>

Journal

Journal of Personality and Social Psychology, 89(4)

ISSN

0022-3514

Authors

Sherman, Jeffrey W
Stroessner, Steven J
Conrey, Frederica R
[et al.](#)

Publication Date

2005-10-01

DOI

10.1037/0022-3514.89.4.607

Peer reviewed

Prejudice and Stereotype Maintenance Processes: Attention, Attribution, and Individuation

Jeffrey W. Sherman
University of California, Davis

Steven J. Stroessner
Barnard College, Columbia University

Frederica R. Conrey
Indiana University

Omar A. Azam
Abn Amro

Three experiments examined the relationship between prejudice and processing of stereotypic information. Higher levels of prejudice were associated with greater attention to and more thorough encoding of stereotype-inconsistent than stereotype-consistent behaviors but only when processing capacity was plentiful (Experiments 1 and 3). High-prejudice participants attributed consistent behaviors to internal factors and inconsistent behaviors to external forces (Experiment 2). Together, these results suggest that high-prejudice people attend carefully to inconsistent behaviors to explain them away but only if they have sufficient resources to do so. Results also showed that low-prejudice but not high-prejudice participants formed individuated impressions by integrating the implications of the target's behaviors (i.e., individuating). High levels of prejudice appear to be associated with biased encoding and judgment processes that may serve to maintain stereotypes.

Keywords: prejudice, stereotyping, individuation, attribution, attention

Whereas *prejudice* is typically defined as the positive or negative evaluations of social groups and their members, *stereotypes* are the knowledge, beliefs, and expectations associated with those groups (e.g., Hamilton & Sherman, 1994; Mackie & Smith, 1998). Though research on these two topics has frequently proceeded in isolation, prejudice and stereotyping are clearly related in important ways (see Mackie & Smith, 1998, for a review). For example, a central theme in social psychological theory has been that stereotyping promotes prejudice and that prejudice reduction depends on stereotype change. Much theorizing also assumes that prejudice should influence the extent of stereotyping. In this regard, prejudice has been shown to moderate stereotype endorsement, such that high-prejudice people typically make more stereotypical judgments than do low-prejudice people (e.g., Devine, 1989; Monteith, 1993). One important theoretical tradition posits that this effect emerges because prejudice encourages the use of stereotypes as a means to justify societal inequality (e.g., Allport, 1954; Katz &

Braly, 1933; Lippmann, 1922). This perspective has received increasing attention in recent years, including empirical support (e.g., Crandall & Eshleman, 2003; Jost & Hunyady, 2002; Jost & Major, 2001; Kay & Jost, 2003; Yzerbyt & Anouk, 2001; Yzerbyt, Rocher, & Schadron, 1997).

However, considering the importance of the relationship between prejudice and stereotyping, there is relatively little direct evidence regarding the impact of prejudice on stereotyping. Research that examines the specific mechanisms by which prejudice may influence stereotyping has been particularly scant. In what ways might prejudice influence social-cognitive processes to maintain stereotypes? To date, most of the research investigating this question has focused on stereotype inhibition processes. This research has shown that for many low-prejudice individuals, stereotyping is deemed to be unacceptable, and stereotype use often is accompanied by feelings of guilt or compunction. As a result, low-prejudice persons are much more likely to actively inhibit stereotypic thinking and responding than are high-prejudice persons (e.g., Monteith, 1993; Monteith, Sherman, & Devine, 1998).

Though clearly important, inhibitory processes may be but one of many routes through which prejudice moderates stereotyping. Specifically, because they regulate the expression but not necessarily the content of stereotypical beliefs, inhibitory processes can be viewed as having an indirect role in stereotyping. It is also possible that prejudice has a more direct role by influencing the content of stereotypical beliefs through cognitive mechanisms employed during the processing of stereotype-relevant information. Prejudice, for example, might affect allocation of attention to information that confirms or contradicts existing stereotypes, might influence how that information is interpreted, and might

Jeffrey W. Sherman, Department of Psychology, University of California, Davis; Steven J. Stroessner, Department of Psychology, Barnard College, Columbia University; Frederica R. Conrey, Department of Psychology, Indiana University; Omar A. Azam, Abn Amro, Chicago, Illinois.

This research was supported by National Institute of Mental Health Grants 55037 and 59774 to Jeffrey W. Sherman and National Science Foundation Grants 9709797 and 9911004 to Steven J. Stroessner. We thank Galen Bodenhausen and Jim Sherman for their helpful comments on an earlier version of this article.

Correspondence concerning this article should be addressed to Jeffrey W. Sherman, Department of Psychology, University of California, One Shields Avenue, Davis, CA 95616. E-mail: jsherman@ucdavis.edu

moderate whether available information is used as the basis for judgments of groups and their members. These relatively direct effects of prejudice on stereotyping are the focus of this research.

In our research, we examine the influence of prejudice on the encoding and integration of behavioral information. Regarding encoding, we investigate the thoroughness with which stereotype-consistent and stereotype-inconsistent behaviors are encoded (Experiments 1 and 3) and the manner in which these behaviors are explained (Experiment 2). We also examine the role of cognitive resources in these processes (Experiment 3). Regarding behavioral integration, we examine the extent to which high- and low-prejudice persons' judgments of a target are based on the target's individuating behaviors (Experiments 1 and 2). Although numerous cognitive processes are implicated in stereotyping, we focus on these processes because they have proven to be central to perceiving individuals and groups (Hamilton & Sherman, 1994) and because of their potential importance for modifying the content of stereotypical beliefs.

Attention/Encoding Quality

One important way in which prejudice may influence stereotyping is by affecting the amount of attention paid to stereotypic and counterstereotypic behaviors and, consequently, the extent to which those behaviors are carefully encoded. If, for example, high-prejudice persons attend to and encode consistent behaviors more thoroughly than inconsistent behaviors, this may contribute to the formation of stereotypic judgments. Conversely, if low-prejudice persons attend to and encode inconsistent behaviors more thoroughly than consistent behaviors, nonstereotypic judgments will likely form.

There are reasons to expect such encoding effects. First, there is a long history of research on selective attention attesting to the fact that people often attend to information that is consistent or congenial with existing attitudes and beliefs (see Frey, 1986, for a review). If consistent behaviors are more congenial to high-prejudice people and inconsistent behaviors are more congenial to low-prejudice people, then one would expect encoding biases consistent with those attitudes. Second, differences in their concerns about forming biased impressions may cause low- and high-prejudice individuals to differentially encode consistent and inconsistent information. Attention tends to be directed at inconsistent rather than consistent information when perceivers are motivated to be accurate, whereas perceivers unconcerned with accuracy are more likely to focus on consistent information (for a review, see Fiske, Lin, & Neuberg, 1999). Given that low-prejudice persons often have strong internal motivations to form nonbiased impressions (Plant & Devine, 1998), they might be expected to more carefully encode inconsistent than consistent information. In contrast, high-prejudice persons, who are not so concerned with accuracy, may be expected to more carefully encode consistent than inconsistent information.

Despite these reasons favoring the encoding of congenial information, there are theoretical bases for predicting other outcomes. If, as argued above, low-prejudice individuals are motivated to form accurate impressions, it is unclear why this motivation would produce a bias favoring inconsistent information. A motivation to be accurate might instead produce equal encoding effort for all

types of available information, without regard to whether the information was consistent or inconsistent with a preexisting expectancy. This would occur if these individuals gave no weight to expectancies (e.g., Brewer, 1988).

There is even reason to predict the opposite pattern to that described above, with high-prejudice persons encoding inconsistent information more carefully and low-prejudice persons encoding consistent information more carefully. In their quantity of processing view of motivated reasoning, Ditto and his colleagues argue that people tend to accept the validity of information that is consistent with desired beliefs with little scrutiny. In contrast, people set a much higher threshold for accepting preference-inconsistent information and, therefore, scrutinize that information very carefully (Ditto & Lopez, 1992; Ditto, Munro, Apanovitch, Scepansky, & Lockhart, 2003). To the extent that high- and low-prejudice people prefer to view others in stereotypic or counterstereotypic terms, respectively, then this may affect the extent to which they process stereotype-consistent and stereotype-inconsistent information. In particular, because they are less motivated to form accurate impressions, more motivated to uphold their stereotypes, and more likely to feel threatened by stereotype-inconsistent information, high-prejudice individuals may attend very closely to this information. In this case, however, effort would not be directed at the inconsistent information to support unbiased, individuated impression formation but rather to scrutinize the information for means by which its impact may be minimized or explained away, thereby bolstering the stereotype. At the same time, stereotype-confirming information may receive relatively little scrutiny, as it confirms what high-prejudice individuals want to believe. For the same reasons, low-prejudice persons might attend carefully to stereotype-confirming information to explain it away in pursuit of forming nonstereotypic judgments and might give little scrutiny to stereotype-disconfirming information.

Ditto and his colleagues have provided evidence for such motivated processes in a variety of domains, including impressions formed of liked and disliked others and self-diagnoses of threatening medical conditions (Ditto & Lopez, 1992; Ditto et al., 2003). In the stereotyping domain, research has shown that people tend to recall more stereotype-inconsistent than stereotype-consistent information about out-group but not in-group members (Bardach & Park, 1996; Wänke & Wyer, 1996), demonstrating that those presumably least motivated to form accurate and/or nonstereotypical impressions (out-group members) may be particularly likely to notice and remember counterstereotypical target information, perhaps because of the careful scrutiny given to this information. In other research, sexist individuals were shown to experience more agitation-related emotions in response to stereotype-inconsistent information, and to have superior memory for that information, than individuals low in sexism (Förster, Higgins, & Strack, 2000). Finally, Eagly and her colleagues have demonstrated that perceivers often attend more carefully to counterattitudinal than proattitudinal messages as they attempt to diminish the impact of uncongenial information (Eagly, Kulesa, Brannon, Shaw, and Hutson-Comeaux, 2000).

Thus, there is precedent to expect that high levels of prejudice may be associated with preferential encoding of either stereotype-consistent or stereotype-inconsistent information. Low levels of prejudice may be associated with preferential encoding of either

inconsistent or consistent information or, alternatively, associated with equivalent encoding of the two types of information. Experiments 1 and 3 were designed to address these competing possibilities.

Attributional Processes

Another important encoding process that may be affected by prejudice is the manner in which stereotypical and counterstereotypical behaviors are explained. If high-prejudice individuals are relatively likely to attribute consistent behaviors to internal causes and inconsistent behaviors to external causes, whereas low-prejudice individuals are relatively likely to attribute consistent behaviors to external causes and inconsistent behaviors to internal causes, then this may explain judgmental disparities between high- and low-prejudice people. This possibility is discussed in greater detail in the introduction to Experiment 2.

Behavioral Integration

A third important mechanism by which prejudice may influence stereotyping is by influencing the extent to which people base their impressions of a person on that person's specific behaviors. If low-prejudice people are more likely to integrate a target's specific behaviors into an impression, whereas high-prejudice people are more likely to base their impressions on prior expectancies, then this also may explain why low-prejudice people form less stereotypical impressions than high-prejudice people. Certainly, to the extent that behavioral integration is a relatively effortful impression formation strategy, we might expect that accuracy-motivated low-prejudice people would be more likely than high-prejudice people to engage in such processes (Brewer, 1988; Fiske & Neuberg, 1990). However, it also may be the case that high-prejudice people are motivated to form behavior-based impressions that are based on only a subset of stereotypical target behaviors. That is, high-prejudice individuals may very carefully integrate a target's stereotype-confirming behaviors into a stereotypical impression while ignoring the person's counterstereotypic actions (e.g., Ditto & Lopez, 1992; Kunda, 1990). These possibilities were investigated in Experiments 1 and 2.

Experiment 1

Overview

Experiment 1 examined the influence of prejudice on the extent to which stereotypical and counterstereotypical information are attended to and encoded. It also assessed the influence of prejudice on the extent to which impressions are based on behavioral integration. High- and low-prejudice participants read stereotype-consistent and stereotype-inconsistent information about a gay man. To assess encoding quality, recognition memory for this material later was assessed. Recognition memory measures the extent to which information has been attended to and encoded well enough to discriminate it from information that has not been encountered (Graesser, 1981; Grier, 1971; Sherman & Frost, 2000; Srull, 1984). Participants also made trait judgments about the target person. Correlations between these judgments and relative recognition accuracy for consistent versus inconsistent information

were used to assess the extent to which judgments were based on the behavioral information that was most thoroughly encoded.

Method

Participants. Fifty-nine undergraduates at Northwestern University participated in the experiment in exchange for partial credit in an introductory psychology course. Participants were recruited on the basis of their responses to the 20-item Heterosexual Attitudes Toward Homosexuals (HATH) scale (Larsen, Reed, & Hoffman, 1980), which had been administered in a mass testing at the beginning of the term. We focused on attitudes toward homosexuals because recent work suggests that holding negative attitudes toward this group is still considered socially acceptable by many (Herek, 1998), allowing a simpler and more explicit assessment of individuals' attitudes than would be the case for many other target groups. We computed HATH scores for each participant by summing responses across all items, reverse scoring where appropriate. HATH scores can range from 20 to 100, with larger scores indicating higher levels of antigay prejudice. Participants with the highest ($M = 72.10$, $n = 21$) and lowest ($M = 21.50$, $n = 22$) HATH scores among all students completing the scale were asked to participate in the study.¹ An additional set of participants was randomly selected from the remaining participants who completed the HATH scale. These participants ($M = 33.00$, $n = 16$) possessed, on average, a moderate level of prejudice. Participants were not told of the true basis for their inclusion in the study and were led to believe they had been randomly selected for participation.

Procedure. Participants were welcomed by an experimenter who was blind to their level of prejudice and were told that the study assessed how people form first impressions of others. They were informed that a number of residents of the Chicago area had been paid to complete interviews regarding details about their lives. Some information gathered during the interviews had been entered into a computer, and the participants were going to have the opportunity to read some of this information later in the experiment. Participants were told that although many individuals had been interviewed, they would have time to read about only one person. To determine which person they would read about, participants selected a piece of paper containing a person's name from a bag containing many small slips of paper. The participant then informed the experimenter that he or she had selected a piece of paper with the name "Robert" written on it. Actually, the name Robert appeared on all of the slips. The experimenter then left the room for several moments, ostensibly to prepare the computer to present information about Robert.

Participants then were seated before a computer monitor. They were informed that Robert, the man they would read about, was a 35-year-old gay man who lives on the north side of Chicago. They were to begin the computer program by pressing a key on the keyboard, and the computer would then present some details about Robert's life. Participants were asked to form an impression of Robert while they read the information.

Participants were then presented with 24 pieces of information about the target, selected on the basis of pilot testing. Eight of the items were consistent with the stereotype of gay men as effeminate (e.g., "studied interpretive dance in college"), eight were inconsistent with the stereotype

¹ Given that approximately 400 students participated in the group testing session, the high- and low-prejudice participants in the experiment represented approximately the top and bottom 5% of respondents. The same is true in Experiment 2. In both experiments, there were equal numbers of male and female subjects selected at the different levels of prejudice. Thus, in Experiment 2, the highest 14 men and 14 women were selected, as were the lowest 14 men and 14 women. Unfortunately, participants' gender was not recorded with their responses. As such, gender could not be included as a factor in the analyses. In Experiment 3, gender was not recorded at all.

(e.g., “watched a football game on Sunday afternoon”), and eight were filler items, irrelevant to the stereotype (e.g., “looked up the telephone number”). In a pilot test, the consistent and inconsistent items were evaluated to be equally positive. The items appeared in a different random order for each participant, and each item appeared on the computer screen for 6 s.

After reading the items on the computer, participants completed a filler task for 5 min. The filler task, which required participants to solve a series of mental puzzles that involved selecting the correct letter in a series of letters (e.g., “What letter comes next in the sequence a, c, e, g?”), was used to clear short-term memory. Prior to completing the recognition task, participants were given the opportunity to familiarize themselves with using the “yes” and “no” keys on their computer keyboards by answering some simple questions (e.g., “Is Coke better than Pepsi?”). The key labeled “yes” was always on the side of the keyboard that corresponded to the participant’s dominant hand.

After completing the practice trials, participants were told that they would be presented with some statements that referred to the target, Robert. Whereas some items reflected information that described Robert in the impression formation phase of the experiment, participants were told, other items were not part of the original description. They were instructed to press the “yes” key if they recognized the item from the original description of Robert and to press the “no” key if they believed that the item had not been part of the description. They were instructed to answer both as quickly and accurately as possible, but they were told not to sacrifice accuracy for speed.

Forty-eight items then appeared on the computer screen in random order. Sixteen of the items were consistent with the stereotype of gay men. Of those items, half had been part of the description of Robert (i.e., “old” items), and half were foils that were not part of the description (i.e., “new” items). Sixteen of the items were inconsistent with the stereotype of gay men. Of those, half were old and half were new. The remaining 16 items, half old and half new, were irrelevant to the stereotype.

Participants then completed a trait-rating task. They were asked to rate the target on several attributes using a 6-point scale (1 = *not at all descriptive*, 6 = *very descriptive*). Embedded among those attributes were two stereotype-relevant traits (feminine and rugged).

After completing the procedure, participants were thanked for their participation and were partially debriefed. Although participants were told that the study pertained to memory for stereotypical information, they were not informed that their inclusion in the study was based on their level of prejudice toward homosexuals.

Results

Trait ratings. Participant’s ratings of the target on the traits feminine and rugged (reverse scored) were submitted to a 3 (prejudice: low vs. medium vs. high) \times 2 (trait) analysis of variance (ANOVA), with repeated measures on the last factor. This analysis yielded a marginal prejudice main effect, $F(2, 56) = 3.00, p = .058$, unqualified by a Prejudice \times Trait interaction, $F(2, 56) = 0.11$, indicating that the effect of prejudice was consistent across traits. Post hoc tests comparing the average ratings showed that participants high in prejudice ($M = 3.74$) rated the target significantly higher on stereotypical traits than did participants low in prejudice ($M = 3.20$), $F(1, 41) = 4.99, p < .05$. Participants with moderate levels of prejudice ($M = 3.56$) rated the target marginally higher on stereotypical traits than did participants low in prejudice, $F(1, 36) = 3.48, p < .08$. Ratings of participants who were moderate and high in prejudice did not differ significantly ($p = .50$). Thus, as expected, participants who were high in prejudice made more stereotypical ratings of the target than did

participants who were low in prejudice. Finally, whereas the ratings of participants high, $t(20) = 1.17, p = .25$, and moderate, $t(15) = 1.17, p = .65$, in prejudice did not differ from the midpoint of the scale, low-prejudice participants rated the target below the midpoint of stereotypicality, $t(21) = -2.27, p < .05$, perhaps reflecting these participants’ responsiveness to the relatively large number of counterstereotypic behaviors.

Recognition for stereotype-relevant information. The proportions of hits (correct identification of old items) and false alarms (failure to reject new items) were used to compute separate measures of recognition accuracy for stereotype-consistent and stereotype-inconsistent items for each participant. A' , a nonparametric measure of recognition sensitivity, was chosen because some participants exhibited perfect memory discrimination. Like other sensitivity measures, A' reflects the degree to which the participant was able to correctly discriminate old items from new items, while controlling for guessing strategies and response biases. The formula for A' is as follows: $A' = .5 + \{[\text{hits} - \text{false alarms}]/[4(\text{hits})(1 - \text{false alarms})]\}$.

Participants’ A' measures were submitted to a 3 (prejudice) \times 2 (behavior stereotypicality: consistent vs. inconsistent) ANOVA, with repeated measures on the second factor. Despite a very high overall recognition rate, this analysis yielded a marginal Prejudice \times Stereotypicality interaction, $F(2, 56) = 2.36, p = .10$ (see Figure 1). To better understand the nature of the interaction, we conducted simple effect analyses at each level of prejudice. For those participants who were high in prejudice, discrimination for stereotype-inconsistent information ($M = .98$) was significantly better than for stereotype-consistent information ($M = .96$), $F(1, 56) = 7.78, p < .01$. In contrast, the same difference was neither significant for participants who were low in prejudice (inconsistent $M = .99$, consistent $M = .99$) nor significant for those who




Figure 1. Recognition sensitivity as a function of prejudice level and behavior stereotypicality in Experiment 1.

exhibited moderate levels of prejudice (inconsistent $M = .99$, consistent $M = .99$), both $F_s < 1$.^{2,3}

The relationship between recognition and judgment. To assess the relationship between encoding and judgment, we derived an index of differential recognition for stereotypical information by subtracting each participant's A' measure for inconsistent information from each participant's A' for consistent information. Therefore, higher scores reflect better recognition of consistent relative to inconsistent information. This index was then correlated with participants' average rating of the target on the stereotypic traits. Despite the rather restricted range of recognition performance, participants who were low in prejudice demonstrated a strong positive correlation between recognition performance and trait ratings, $r(21) = .44$, $p < .05$. Thus, the tendency to make more stereotypical responses was associated with better recognition of stereotype-consistent information. For participants who were moderate in prejudice, recognition was uncorrelated with responses to the trait measure, $r(16) = -.05$. Finally, for participants who were high in prejudice, there was a nonreliable negative correlation between recognition for consistent information and judgment stereotypicality, $r(20) = -.33$, $p = .16$. The difference between the low- and high-prejudice participants' correlations was reliable, ($p < .05$).

Discussion

The data from Experiment 1 demonstrate that level of prejudice influences the encoding and judgment of stereotype-prone targets. As was expected, high-prejudice participants made more stereotypical judgments of a gay target than did low-prejudice participants. At the same time, the recognition data showed that high-prejudice persons encoded stereotype-inconsistent information more thoroughly than stereotype-consistent information. There are two possible interpretations of the differences in encoding of stereotype-relevant information. One possibility is that high-prejudice participants focused on processing inconsistent information, perhaps as they attempted to explain away the unexpected behaviors. Second, the effect might indicate relatively minimal processing of consistent information that confirms stereotypic expectancies and, as a result, may stimulate particularly little explanatory effort among these high-prejudice participants. Indeed, the pattern of results suggests that diminished processing of consistent items may have played a larger role in the effect than did enhanced processing of inconsistent items. Of course, both enhanced processing of inconsistent and diminished processing of consistent items might have contributed to the recognition advantage for inconsistent items, a possibility that is examined in Experiments 2 and 3. Nonetheless, it is very clear that high-prejudice participants' attention was not biased toward stereotype-confirming information, and low-prejudice participants' attention was not biased in a stereotype-disconfirming fashion. These findings are consistent with the suggestion that people set a higher threshold for accepting undesired than desired information (e.g., Ditto & Lopez, 1992; Ditto et al., 2003; Kunda, 1990).

The correlational data assessing the relationship between recognition and judgment demonstrate that prejudice also influences the manner in which judgments are made. Participants who were low in prejudice appear to have integrated the behavioral information

into individuated impressions of the target. To the degree they exhibited a pattern of stereotype-confirming (disconfirming) encoding, judgment stereotypicality increased (decreased). In contrast, high-prejudice participants showed no relationship between recognition and judgment. This suggests that high-prejudice participants were not forming systematic, individuated impressions on the basis of the behaviors they attended to and encoded. Rather, it seems likely that they were basing their judgments on preexisting stereotypes. Thus, whereas high-prejudice participants were more carefully encoding inconsistent than consistent information but were not using the behaviors as a basis for their impressions, low-prejudice participants were encoding consistent and inconsistent information with equal care and were using the implications of whichever behaviors caught their attention to inform their impressions.⁴

Experiment 2

Prejudice and Attribution

The results of Experiment 1 suggest that high-prejudice persons more thoroughly encode stereotype-inconsistent than stereotype-

² On the basis of the responses to the recognition test, nonparametric B' indices of response bias toward consistent and inconsistent items also were computed (Grier, 1971). These indices were analyzed in a 3 (prejudice) \times 2 (behavior stereotypicality: consistent vs. inconsistent) ANOVA, with repeated measures on the last factor. This analysis produced only a main effect for behavior stereotypicality such that response bias was stronger for consistent ($M = .11$) than inconsistent ($M = .02$) items, $F(1, 56) = 7.50$, $p < .01$, replicating a common finding that perceivers set a lower criterion for responding positively to expectancy-consistent than expectancy-inconsistent information (Graesser, 1981; Sherman & Frost, 2000; Srull, 1984). This effect was not moderated by level of prejudice.

³ Because of the restricted range of the HATH scores within each prejudice group, regression analyses are not appropriate for Experiments 1 or 2. This is particularly the case for low-prejudice participants, who had a mean on HATH of 21.5, with 20 being the lowest score possible.

⁴ An alternate explanation of these results is that the high-prejudice participants possessed stronger stereotypes than did the low-prejudice participants. Such differences in stereotype strength might lead high-prejudice participants to make stereotypic judgments and to notice and attend to inconsistencies because those items are particularly surprising or difficult to comprehend (Sherman et al., 1998; Srull, Lichtenstein, & Rothbart, 1985). However, extensive pilot testing on the relationship between attitudes and expectancies about gay men suggests this is unlikely to have been the case. A separate group of participants ($n = 24$) was asked to complete the HATH scale, a feeling thermometer (an alternative attitude measure; Haddock, Zanna, & Esses, 1993), and a trait rating scale. The rating scale prompted participants to rate homosexual men and people in general on stereotype-relevant traits (e.g., emotional, artistic, weak, feminine) and to give their confidence in these ratings. These ratings were then used to compute metacontrast ratios (the extent to which traits were rated as more typical of homosexuals than people in general), reflecting the strength of participants' gay stereotype. Scores from the HATH scale were highly correlated with ratings on the feeling thermometer ($r = -.46$, $p < .05$) but uncorrelated with measures of stereotype strength ($r = .03$). Moreover, participants' confidence in their trait ascriptions was also uncorrelated with responses on the HATH scale ($r = -.08$). These data reinforce the status of the HATH scale as an attitudinal measure, and suggest that our results cannot simply be attributed to stereotype strength.

consistent information. One possible factor in this result is the kinds of explanations high- and low-prejudiced people generate for consistent and inconsistent behaviors. As described above, high-prejudiced people may set a high threshold for accepting stereotype-disconfirming information and, thus, may be highly motivated to explain away inconsistent behaviors. If a behavior can be attributed to a situational factor, then the behavior is not necessarily at odds with the group stereotype, and the validity of that stereotype can be maintained. At the same time, high-prejudice people may set a low threshold for accepting stereotype-confirming information and, thus, may be willing to take consistent behaviors at face value as reflections of the true disposition of the actor that require no further explanation (e.g., Ditto & Lopez, 1992; Ditto et al., 2003). The greater effort expended on generating situational explanations for inconsistent but not consistent behaviors may contribute to high-prejudice participants' relatively thorough encoding of and memory for inconsistent behaviors (Hastie, 1984).

In contrast, to the extent that low-prejudiced participants are interested in disconfirming stereotypes, they may be more likely to make internal attributions for inconsistent behaviors and external attributions for consistent behaviors. Alternatively, if low-prejudiced participants are interested primarily in forming unbiased impressions, we would expect that internal and external attributions would be equally common for consistent and inconsistent behaviors. Consistent with this possibility, Experiment 1 provided no indication that consistent and inconsistent items were treated differently by low-prejudice persons.

Biases in attributional processes have been demonstrated in a number of lines of research but have not been clearly delineated in the literature on prejudice. Erber and Fiske (1984), for example, demonstrated that participants who were unmotivated to form accurate impressions made dispositional attributions for consistent behaviors and situational attributions for inconsistent behaviors. In contrast, those who were motivated via outcome dependency to form accurate impressions made more situational attributions for consistent behaviors than for inconsistent behaviors and an equal number of dispositional attributions for consistent and inconsistent behaviors. In other research, von Hippel, Sekaquaptewa, and Vargas (1997) showed that higher levels of prejudice were associated with an increased likelihood of generating causal explanations for inconsistent but not consistent behaviors. However, von Hippel et al. did not distinguish between internal and external attributions in their analyses. In Experiment 2, we directly examine the extent to which high- and low-prejudice participants generate internal and external attributions for consistent and inconsistent behaviors.

Behavioral Integration

The second purpose of Experiment 2 was to examine in more detail the issue of prejudice and behavioral integration processes. Experiment 1 showed that for low-prejudice participants, recognition performance was strongly correlated with judgment stereotypicality, indicating that these participants were basing their judgments on the information to which they attended (i.e., they were individuating). In contrast, high-prejudice participants showed no correlation between recognition and judgment stereotypicality, indicating that these participants were not basing their judgments on the information they encoded but on other information (pre-

sumably, stereotypes). However, these data are far from conclusive. First, because the data are correlational, the direction of causality is unclear. Among the low-prejudice participants, it could be that the online formation of relatively stereotypical or nonstereotypical judgments (through whatever means) caused attention to shift toward consistent or inconsistent behaviors, respectively. In this case, the correlation would not necessarily reflect behavioral integration processes.

Second, the correlations between recognition and judgment are a rather crude means of assessing the extent of behavior integration. For example, if the high-prejudice participants in Experiment 1 were engaging in the sorts of attributional biases described above, then we would not expect to see a correlation between recognition and judgment, even if behavioral integration was occurring. Working hard to generate external attributions for inconsistent behaviors would increase subsequent recognition of those items. At the same time, generating external attributions would diminish the extent to which those items disconfirm the stereotype, thereby enhancing judgment stereotypicality. This would mitigate the likelihood of observing a correlation between recognition and judgment and raise doubts about the meaning of any correlation, or lack thereof. In Experiment 2, we more directly assess the extent to which participants' judgments reflect the behavioral evidence.

Overview

As in Experiment 1, high- and low-prejudice participants read stereotype-consistent and stereotype-inconsistent information about a gay man. However, in this experiment, the proportion of consistent and inconsistent behavior was manipulated between participants. Whereas some participants read about a target that performed twice as many consistent as inconsistent behaviors, others read about a target that performed twice as many inconsistent as consistent behaviors. If, on the one hand, participants are forming impressions on the basis of the behavioral evidence, then the generally stereotypical target should be judged to be more stereotypical than the generally counterstereotypical target. If, on the other hand, the two targets are judged to be highly similar, then participants clearly are not relying on the behavioral evidence. The logic of this experiment is analogous to studies of persuasion that measure responses to strong and weak persuasive messages. In that research, participants are presumed to be processing systematically to the extent that they are more persuaded by strong than weak arguments (for reviews, see Eagly & Chaiken, 1993; Petty & Wegener, 1998).

To examine differences in the types of attributions high- and low-prejudice participants make for consistent and inconsistent behaviors, we used a sentence-completion paradigm developed by Hastie (1984). After reading about the target and judging his traits, participants were again presented with the descriptions of his behavior and were asked to extend or embellish them in some meaningful way. These sentence continuations were then coded to identify the proportion of internal or external attributions that were made to account for the behaviors.

Method

Participants. Participants were the 56 Northwestern University undergraduates with the highest ($n = 28$) and lowest ($n = 28$) scores on the

HATH scale collected in an initial screening of introductory psychology students at the beginning of the academic term. These students were solicited to participate in the experiment without being made aware of the basis for their inclusion in the study. All participants received partial course credit for completing the experiment.

Procedure. The procedures for Experiment 2 were similar to those used in Experiment 1. Although students were told that they would be asked to form an impression of a person randomly selected from a pool of individuals, all participants read 24 descriptions of behaviors purportedly performed by Robert, a gay man from Chicago. The array of behavioral descriptions presented, however, differed from Experiment 1. In Experiment 2, all participants received 12 descriptions of behaviors that were irrelevant to the stereotype of gay men. Whether the remaining items were predominantly consistent or inconsistent with the stereotype was manipulated between conditions. In the stereotypical target condition, Robert performed eight stereotype-consistent and four stereotype-inconsistent behaviors. By contrast, in the counterstereotypical target condition, Robert performed four consistent and eight inconsistent behaviors. As in Experiment 1, these behaviors appeared on a computer screen in random order for 6 s each.

After reading the behavioral descriptions, participants completed the same trait-rating task (i.e., feminine and rugged) as in Experiment 1. They were then provided with the behavioral descriptions they had read in the impression formation task in one of two random orders. As in Hastie's (1984) study, participants were asked to reread each behavioral description and write a completion for that sentence. A completion was defined for participants as an addition to the original item that provides a plausible and meaningful extension to the original description. It was suggested that multiple completions were possible for each item and that there was no one correct way to do the task. They were given 30 s to generate and write a completion for each behavioral description. Finally, participants were thanked and debriefed.

Results

Trait ratings. Trait ratings were analyzed by use of a 2 (prejudice) \times 2 (target stereotypicality) \times 2 (order) \times 2 (trait) ANOVA, with repeated measures on the last factor. This analysis produced a significant main effect of prejudice, $F(1, 48) = 10.30$, $p = .05$, showing that, overall, high-prejudice participants made more stereotypical judgments than did low-prejudice participants. However, this effect was moderated by target stereotypicality, producing a significant interaction (see Figure 2), $F(1, 48) = 4.02$,




Figure 2. Mean judgments on stereotypical traits as a function of prejudice level and target stereotypicality in Experiment 2. Higher trait ratings reflect more stereotypical judgments.

$p = .05$. Separate one-way ANOVAs revealed that low-prejudice participants made more stereotypical trait ratings of targets who were stereotypical ($M = 3.39$) than targets who were counterstereotypical ($M = 2.68$), $F(1, 24) = 6.26$, $p < .05$. In contrast, high-prejudice participants' ratings of the two targets did not differ ($M_s = 3.64$ and 3.74). Low and high prejudice participants' ratings of the counterstereotypical target were significantly different from one another, $F(1, 28) = 14.29$, $p < .05$. However, the two groups' ratings of the stereotypical target did not differ, $F(1, 24) = 1.21$, $p = .28$. For neither target did high-prejudice participants' ratings differ from the midpoint of the scale: stereotypical target, $t(12) = 0.84$, $p = .42$; counterstereotypical target, $t(14) = 0.89$, $p = .39$. In contrast, though low-prejudice participants' ratings of the stereotypical target did not differ from the midpoint, $t(12) = -0.72$, $p = .49$, their ratings of the counterstereotypical target were significantly lower than the midpoint, $t(14) = -4.58$, $p < .01$. These data demonstrate that participants who were low, but not those who were high, in prejudice were sensitive to the nature of the target's behavior, and the data also support the contention that only low-prejudice participants were engaged in behavioral integration.

Sentence completions. Two raters who were blind to the participants' prejudice level coded completions. Each completion to a behavioral sentence (e.g., "put in an Ozzy Osborne compact disc") was coded as reflecting either an internal attribution (e.g., "because the music moved him emotionally"), an external attribution (e.g., "to make his younger brother happy"), a noncausal continuation (e.g., "and read a magazine article he was interested in"), or a meaning reversal (e.g., "to spoof Ozzy's macho male posturing"). Agreement between the two coders was high (72% across all categories and 71% for attributions), and disagreements were resolved through discussion with Jeffrey W. Sherman. Overall, 27% of the completions constituted internal attributions, 33% were external attributions, 39% were continuations, and less than 1% of completions constituted reversals.

Because predictions focused on the conditions under which participants would make internal versus external attributions for behavior, only the attributional responses were entered into a 2 (prejudice) \times 2 (target stereotypicality) \times 2 (order) \times 2 (behavior stereotypicality) \times 2 (attribution type: internal vs. external) ANOVA, with repeated measures on the last two factors. This analysis yielded main effects for prejudice, $F(1, 48) = 4.90$, $p < .05$, indicating that participants high in prejudice made a higher proportion of attributions ($M = 0.34$, collapsed across both types of attributions) than did participants low in prejudice ($M = 0.26$), and for behavior stereotypicality, $F(1, 48) = 5.06$, $p < .05$, replicating the common finding that unexpected behaviors ($M = 0.32$, collapsed across both types of attributions) prompt more attributional processing than do expected behaviors ($M = 0.28$) (Hastie, 1984). A Behavior Stereotypicality \times Attribution Type interaction also emerged, $F(1, 48) = 12.13$, $p < .001$, reflecting the fact that there were more external ($M = 0.40$) than internal ($M = 0.23$) attributions for stereotype-inconsistent behaviors, $F(1, 55) = 10.41$, $p < .01$, but no difference in the types of attributions made for stereotype-consistent behaviors ($M_s = 0.26$ and 0.30 , respectively).

These effects were qualified by a significant three-way interaction involving prejudice, behavior stereotypicality, and attribution type, $F(1, 48) = 4.04$, $p = .05$. To better understand this interac-

tion, we conducted separate 2 (behavioral stereotypicality) \times 2 (attribution type) analyses for participants who were low (see Figure 3, left panel) and high (see Figure 3, right panel) in prejudice. For the low-prejudice participants, no effects were significant, $ps > .29$, indicating that the numbers of internal and external attributions for consistent and inconsistent behaviors did not differ. In contrast, the analysis for participants high in prejudice yielded a significant interaction, $F(1, 27) = 22.94, p < .001$. Whereas more internal attributions were made for consistent ($M = .36$) than inconsistent ($M = .22$) behaviors, $F(1, 27) = 8.20, p < .01$, more external attributions were made for inconsistent ($M = .49$) than consistent ($M = .28$) behaviors, $F(1, 27) = 37.96, p < .001$. This demonstrates that high-prejudice participants indeed were engaged in biased attributional processing, seeking to explain away the inconsistent behaviors but accepting the consistent behaviors as reflections of the target's attributes.

Discussion

Experiment 2 extends the results from Experiment 1 in a number of important ways. First, the relationship between prejudice and encoding processes is clarified. As expected, high- but not low-prejudice participants were more likely to seek external attributions for inconsistent than consistent behaviors. It appears that one reason that inconsistent behaviors attract attention from high-prejudice people is that they are motivated to explain them away. At the same time, high- but not low-prejudice participants were more likely to make internal attributions for consistent than inconsistent behaviors. This may explain why consistent behaviors receive relatively little attention from high-prejudice people—these behaviors are simply accepted at face value (e.g., Ditto & Lopez, 1992; Ditto et al., 2003). Thus, whereas the results of Experiment 1 suggested that high prejudice participants' bias was due largely to a decrease in the extent to which they encoded consistent information, the results of Experiment 2 demonstrated that high-prejudiced people process both consistent and inconsistent information in a biased fashion.

Regarding low-prejudice participants, these data support the notion that they are largely motivated to seek accurate impressions of members of stereotyped groups. Just as they encoded consistent and inconsistent behaviors equally thoroughly in Experiment 1, so too did they make an equal proportion of internal and external attributions for consistent and inconsistent behaviors in Experiment 2.

Second, Experiment 2 provided more direct evidence regarding behavioral integration processes. Echoing the results from Experiment 1, Experiment 2 showed that low- but not high-prejudice participants engaged in systematic, behavioral integration processes. In particular, the data showed that low- but not high-prejudice participants distinguished between the stereotypical and the counterstereotypical target. Among the low-prejudice participants, when the target was counterstereotypical, he was rated accordingly. Similarly, if the target behaved in a stereotypical manner, he was judged to be relatively stereotypical. Indeed, low-prejudice participants rated the stereotypical target to be as stereotypical as did the high-prejudice participants. This shows that the low-prejudice participants were not simply inhibiting their stereotypes. Rather, they were integrating the actor's behaviors into a truly individuated impression, even if those integration processes resulted in an impression that resembled the stereotype. In contrast, the fact that high-prejudice participants saw the two targets as more or less identical indicates that they were not basing their impressions on the targets' behaviors.

Of course, high-prejudice participants may not have judged the two targets to differ from one another because they had explained away the stereotype-inconsistent behaviors so successfully. In this case, it is possible that these participants' judgments reflected biased integration processes, in which the products of the biased attributions were integrated into a target impression. However, subsequent analyses of the correlations between participants' attributions and their trait judgments demonstrated no relationship between the two. It seems that high-prejudice participants simply did not take the target's behaviors into account in forming their impressions, regardless of how those behaviors were interpreted.




Figure 3. Mean proportion of internal and external attributions for stereotype-consistent and stereotype-inconsistent behaviors for participants low (left panel) and high (right panel) in prejudice in Experiment 2.

Instead, these participants appear to have relied heavily on their stereotypes of gay men in forming their impressions.

Experiment 3

There were three goals in Experiment 3. First, we wanted to replicate the findings of Experiment 1 regarding prejudice and attention/encoding effort using a different group stereotype. Second, we wanted to examine attention with a more fine-grained and direct measure. Finally, we wanted to examine the role of cognitive resources in directing perceivers' attention toward consistent and inconsistent information.

Whereas in Experiment 1 we focused on stereotypes of gay men, in Experiment 3 we focused on stereotypes of African Americans to assess the generality of our findings. However, because prejudice based on race is more sensitive to self-presentational concerns than is prejudice based on sexual orientation, explicit measures (i.e., racism scales) may be less valid measures of attitudes about race than sexual orientation (e.g., Fazio, Jackson, Dunton, & Williams, 1995). Indeed, correlations between explicit and implicit measures of heterosexism are stronger than correlations between explicit and implicit measures of racism (Nosek, 2004). Moreover, whereas explicit antigay prejudice (i.e., on the HATH) is readily observed among Northwestern University students, anti-Black prejudice is rarely evident among these students on explicit measures of racism. As such, in Experiment 3, prejudice against African Americans was assessed with an implicit measure. Specifically, the Go-No Go Association Task (GNAT; Nosek & Banaji, 2001) was used to assess prejudice.

A different measure of attention was used because recognition memory performance in Experiment 1 approached ceiling. Though this ceiling effect may have undermined identification of differences in encoding (which, nevertheless, were clearly evident for high-prejudice participants), such high performance might have completely concealed other important effects both within and between item types and levels of prejudice. Moreover, though recognition memory reflects the extent to which information has been thoroughly encoded and represented in memory, it can reflect other components of encoding and retrieval as well. Beyond attention or effort, other possible contributors to recognition performance are differences in perceptual or conceptual encoding (e.g., Richardson-Klavehn & Bjork, 1988; Sherman, Lee, Bessenoff, & Frost, 1998) and differences in item familiarity or recollection (e.g., Jacoby, Toth, & Yonelinas, 1993). Thus, the recognition data from Experiment 1 are suggestive but not definitive regarding the relationship between prejudice and encoding effort.

Finally, some participants were placed under a cognitive load as they performed the task. Previous research suggests that attention is more likely to be directed toward motivation-congruent information when processing resources are low (Plaks, Stroessner, Dweck, & Sherman, 2001; Sherman et al., 1998). For example, Plaks et al. (2001) showed that people who are motivated to see group stereotypes as malleable (i.e., "incremental theorists") were especially likely to shift attention away from stereotype-consistent and toward stereotype-inconsistent information when they were under cognitive load. In contrast, people who are motivated to see group stereotypes as stable (i.e., "entity theorists") were especially likely to shift attention away from stereotype-inconsistent and

toward stereotype-consistent information when they were under cognitive load. These researchers argued that because resources must be distributed more strategically when capacity is depleted, it is in these conditions that attitude-congruent selective attention effects are most likely to be obtained. Thus, although accuracy-motivated perceivers may generally shift attention away from stereotypic and toward counterstereotypic information when resources are depleted (e.g., Sherman et al., 1998), other processing motives may moderate (and even reverse) this effect.

In terms of the present research, this suggests that the encoding advantage for stereotype-inconsistent information exhibited by high-prejudice individuals may be restricted to conditions in which processing capacity is plentiful. When placed under a cognitive load, this tendency may be eliminated or even reversed, with higher levels of prejudice corresponding to attending more carefully to consistent information, which is presumably more congenial to highly prejudiced attitudes. This may be particularly likely if, as Experiment 2 suggests, high-prejudiced participants' attention is directed at inconsistent information only so that it may be explained away (e.g., Ditto & Lopez, 1992). Such explanatory processes require abundant attention and are significantly curtailed when capacity is diminished (e.g., Gilbert, Pelham, & Krull, 1988; Srull, 1981). If unable to explain away the inconsistent information, high-prejudice participants may find the information threatening and withdraw their attention (Frey, 1986).

Overview

Participants varying in prejudice toward African Americans were asked to form an impression of a Black man under a cognitive load or no load on the basis of descriptions of his behavior. The behavioral descriptions were presented in pairs, with one item on the left side of a computer screen and one item on the right side. At some point in the presentation of each pair, an X appeared on either the left or right side of the screen. Participants' task was to press a key indicating the location of the X as quickly as possible. Among the pairs of behaviors were 10 that included one stereotype-consistent and one stereotype-inconsistent behavior. For these pairs, the X appeared on the side of the screen corresponding to the side on which either the consistent or inconsistent item had appeared. Interest centered on participants' latencies to respond to the X as a function of processing capacity and whether the X appeared in the same position as the consistent or inconsistent behavior. These latencies measure the extent to which participants were attending to the consistent and inconsistent items of a pair when the X appeared. The more attention being paid to a particular item, the less time it should take to respond to an X probe that appears in the same position as that item (e.g., Bradley, Mogg, White, Groom, & de Bono, 1999; Sherman, Conrey, & Groom, 2004).

Method

Participants. Participants were 68 Northwestern University undergraduates who received partial course credit for completing the experiment.

X-probe procedure. Participants were asked to form an impression of a person randomly selected from a pool of individuals from the Chicago area. Ostensibly, participants' computers randomly selected a person, when in fact, all participants read about the same individual. The computer

displayed a picture of the impression target, a young adult Black man wearing a black headband and dark sunglasses. After viewing the photograph for five seconds, participants were introduced to the X-probe task described above. They were given some practice performing the task with pairs of statements unrelated to personality characteristics or stereotypes. After this practice set, participants performed the same task while reading about the behaviors of the target.

The target's behavior was described with 60 sentence fragments that had been pretested for their level of kindness/hostility. Twenty each of the behaviors reflected kindness (e.g., "visited a friend who was sick in the hospital") and hostility (e.g., "insulted the old woman"), and 20 were irrelevant to kindness/hostility (e.g., "bought a new shirt"). Because hostility is a central component of the Black stereotype (e.g., Devine & Baker, 1991), the hostile behaviors were stereotype consistent, and the kind behaviors were stereotype inconsistent. The items were presented on participants' computer screens in pairs, with one on the left and one on the right side. The particular items selected for each pair were generated randomly by the computer without replacement, with the constraint that of the 30 pairs of behaviors, 10 contained consistent and inconsistent items, 10 contained consistent and irrelevant items, and 10 contained inconsistent and irrelevant items. Each pair was presented for 3 s.

At one of four randomly generated times during the presentation of each pair of behaviors, an X appeared on either the left or right side of the computer screens. The X appeared either after 1,500 ms; 1,750 ms; 2,000 ms; or 2,250 ms. Participants were instructed to press keys marked "left" or "right" as quickly as possible to indicate on which side of the screen the X had appeared. The computer recorded the response times.

During the impression formation task, half of the participants were placed under a cognitive load. These participants were informed that the experiment was concerned with people's ability to perform multiple tasks at the same time. Cognitive load was manipulated by asking these participants to hold an eight-digit number in memory as they performed the X-probe task. This manipulation has been used successfully in past research to deprive participants of processing resources (e.g., Sherman & Frost, 2000; Sherman et al., 1998). As a means of assessing compliance, these participants were asked to write down the eight-digit number on a slip of paper at the end of the X-probe task.⁵

Implicit prejudice toward Black men. Upon completion of the X-probe task, participants completed a GNAT to measure their implicit level of prejudice toward Black men. The GNAT measures the strength of association between a target category and two poles of an attribute dimension. In the present experiment, the GNAT measured the strength of association between Black men and positive/negative valence. In one key block of trials, participants were asked to categorize together and respond identically to pictures of Black men and words representing positive concepts (i.e., *rainbow, love, peace*). In another key block of trials, participants were asked to categorize together and respond identically to pictures of Black men and words representing negative concepts (i.e., *murder, hate, death*). The extent to which it is easier to categorize together pictures of Black men and negative concepts than it is to categorize together pictures of Black men and positive concepts reflects an implicit negative evaluation of Black men, that is, implicit prejudice.

The GNAT was introduced to participants as a word-search task aimed at studying the cognitive processes involved in distinguishing categorical stimuli from one another. There was a series of eight practice blocks before the final two blocks of interest. Each practice block consisted of 20 trials. On each trial, a single stimulus item was presented, and participants were told to press the space bar on their computers if the item belonged to a target category of interest and to do nothing if the item did not belong to that category. For all of the practice blocks, participants were given a response deadline of 850 ms, by which time a response (if the item required one) needed to be given. A 300-ms interstimulus interval separated the end of a trial and the beginning of the next trial. A trial ended when the

participant pressed the space bar or when the response deadline was reached, whichever came first. During each block, the target category label remained in the upper part of the computer screens as a reminder.

For the first two practice blocks, the stimuli consisted of 10 positive and 10 negative words, selected randomly and without replacement from pools of 24 positive and 24 negative words. In the first practice block, participants were to press their space bars if a negative word appeared, and positive words were to be treated as distractors, requiring no response. In the second block, participants were to press their space bars if a positive word appeared, and negative words were distractors.

After these two blocks, six other practice blocks were presented in random order. For these blocks, the stimuli were 4 pictures of Black men, Black women, White men, and White women, and 2 pictures each of Asian men and Asian women that were selected randomly and without replacement from a pool of 21 pictures of Black men, Black women, White men, and White women, and 5 pictures each of Asian men and Asian women. Each of the six Race \times Sex categories served as the target category in one of the six practice blocks, with the other five categories serving as distractors. For example, in one case, participants were told to press their space bars if a picture of a White woman appeared and to do nothing if any other kind of picture appeared.

Throughout the GNAT, trials on which the space bar was pressed incorrectly in response to a distractor item (i.e., false alarms) and trials on which a response was withheld incorrectly in response to a target item (i.e., misses) were scored as errors. On these trials, the word *Error* appeared in red below the stimulus item for 100 ms during the interstimulus interval to provide performance feedback. Trials on which the space bar was pressed in response to a target item and trials on which no response was offered to distractor items were noted as correct responses with a green *Correct*.

After completion of the eight practice blocks, the two key blocks were given in random order. The stimuli for these blocks included 12 positive and 12 negative words, eight pictures of Black men, Black women, White men, and White women, and four pictures of Asian men and Asian women, for a total of 64 trials. The stimuli were randomly selected from the pools used for the practice trials, described above. In one of the blocks, participants were instructed to press their space bars only if either a picture of a Black man or a positive word appeared. If any other type of stimulus appeared, they were to do nothing. In the other block, they were instructed to press their space bars only if either a picture of a Black man or a negative word appeared and to do nothing if any other type of stimulus appeared. A response deadline of 500 ms was given for each block. Together these two blocks assess the strength of association between Black men and negativity/positivity compared with the strength of association between the other five Race \times Sex categories and negativity/positivity.

Results

Calculation of implicit prejudice. With the GNAT, implicit prejudice is calculated via response sensitivity (Nosek & Banaji, 2001). Sensitivity indicates the ability to discriminate target items (signal) from distractors (noise), while response bias is controlled for. The assumption is that sensitivity (separating signal from noise) ought to be easier when the two targets are positively associated than when they are not related or are negatively associated. Thus, to the extent that Black men are implicitly associated with negativity and not with positivity, sensitivity will be higher when Black men and negative words are jointly discriminated from distractors than when Black men and positive words are

⁵ Three participants misreported more than four of the numbers. Results that included and excluded these participants were identical. The data reported include those participants.

jointly discriminated from distractors. Greater sensitivity indicates a stronger association between the target category (Black men) and the attribute (negative words vs. positive words).

The proportion of hits (correct responses to target items) and false alarms (incorrect responses to distractors) were used to compute separate A' measures of sensitivity for the pairing of Black men with negative words and for the pairing of Black men with positive words. For example, with the pairing of Black faces and negative words, space bar responses to Black faces and negative items are hits, and responses to any other type of face or positive items are false alarms. These measures of sensitivity were submitted to a 2 (load: low vs. high) \times 2 (association type: Black–bad vs. Black–good) ANOVA, with repeated measures on the second variable. This analysis produced a main effect of association type, $F(1, 66) = 14.90, p < .001$, indicating greater sensitivity for the pairing of Black with bad ($M = .85$) than for the pairing of Black with good ($M = .80$). Thus, overall, participants associated Black men more with negativity than with positivity.

Individual variation in implicit prejudice and attention. We calculated individual estimates of implicit prejudice for each participant by subtracting response sensitivity for pairing Black men with positive words from response sensitivity for pairing Black men with negative words. Thus, higher scores equal greater relative ease of pairing Blacks with negativity than with positivity (i.e., implicit prejudice).

We calculated individual estimates of attention for each participant by first eliminating incorrect responses (5%) and responses 2.5 standard deviations slower than a given participant's mean ($M = 1.39\%$ per participant). We then inversely transformed ($1/x$) the remaining response times to normalize the data and computed an attentional bias index by subtracting the mean inverse latency to respond to probes appearing at the same location as stereotype-inconsistent items from the mean inverse latency to respond to probes appearing with stereotype-consistent items. Thus, higher scores reflect relatively faster responses to Xs appearing with consistent versus inconsistent items (i.e., a bias favoring consistent over inconsistent items).^{6,7}

We conducted a simultaneous regression analysis to examine the influence of implicit prejudice, cognitive load, and their interaction on the relative amount of attention paid to consistent and inconsistent items. The omnibus model was significant, $F(3, 64) = 4.88, p < .01, R^2 = .19$. The main effect for cognitive load was not reliable ($p = .22$). There was a significant main effect for level of prejudice, $\beta = -.25, t(64) = -2.15, p < .05$, indicating that higher levels of prejudice were associated with paying greater attention to inconsistent information. However, this effect was moderated by cognitive load, $F(1, 64) = 5.91, p < .05$ (see Figure 4). Analyses of the interaction demonstrated that the relationship between prejudice and attention to inconsistent information was strong and significant in the low-load condition, $\beta = -.62, t(32) = -4.48, p < .0001$, providing a conceptual replication of the results from Experiment 1, but nonexistent in the high-load condition, $\beta = .03, t(32) = .16, p = .88$.

To analyze these results in greater detail, we examined the attention paid to consistent and inconsistent items separately within each load condition rather than as a difference score. Analyses of responses in the low-cognitive load condition revealed a significant interaction between level of prejudice and item con-

sistency, $F(1, 32) = 20.10, p < .0001$ (see Figure 5, left panel). Closer examination showed a marginal positive correlation between prejudice and attention to inconsistent information, $\beta = .28, t(32) = 1.67, p = .10$, and a nonsignificant negative correlation between prejudice and attention to consistent information, $\beta = -.06, t(32) = -0.34, p = .73$. A separate analysis of the high-cognitive load condition produced no significant effects (see Figure 5, right panel).

Discussion

The results of Experiment 3 replicate and extend the findings of Experiment 1. Using a different stereotype and a more direct measure of attention, Experiment 3 again showed that higher prejudice was associated with paying greater attention to stereotype-inconsistent than stereotype-consistent information. However, this relationship held only when participants had full processing capacity. When placed under a cognitive load, individuals showed no bias in attention as a function of prejudice. The attenuation of this effect in the high load condition is consistent with the idea that high-prejudice persons were attending to inconsistent information primarily to discount it or to explain it away. When deprived of the resources to engage in such attributional efforts, the impetus to attend carefully to inconsistent items was removed.

In Experiment 1, the results suggested that the relationship between prejudice and encoding might have been largely due to a diminishment in the extent to which high-prejudice participants encoded consistent information. In Experiment 2, the results suggested that high-prejudiced people process both consistent and inconsistent information in a biased fashion. In Experiment 3, higher prejudice was associated only with increasing attention to inconsistent items in the low load condition (see Figure 5, left panel). Together, the results of all three experiments suggest that prejudice may be related to biased processing of either consistent or inconsistent information or to both.

General Discussion

This research aimed to specify the influence of prejudice on the encoding and integration of behavioral information. The results showed that those with high levels of prejudice readily engage in processing stereotype-inconsistent information, particularly if they have full processing capacity. In Experiments 1 and 3, higher levels of prejudice were associated with attending to and encoding inconsistent information more thoroughly than consistent information. This finding was obtained with two different social groups and with both an explicit (Experiment 1) and an implicit (Exper-

⁶ Because analyses are based on inverse scores, higher scores equal faster times.

⁷ Whether the X probe appeared on the left or right side of the screen was included as a factor in an initial analysis. This analysis showed that side of the screen did not interact with stereotype consistency or level of prejudice and did not moderate any of the reliable effects. Including this factor in the analysis produces a greater number of empty cells that are due to errors and outliers (because there are fewer total responses within each cell) and, as such, reduces the number of participants available for analysis. Consequently, this factor was dropped from subsequent analyses.

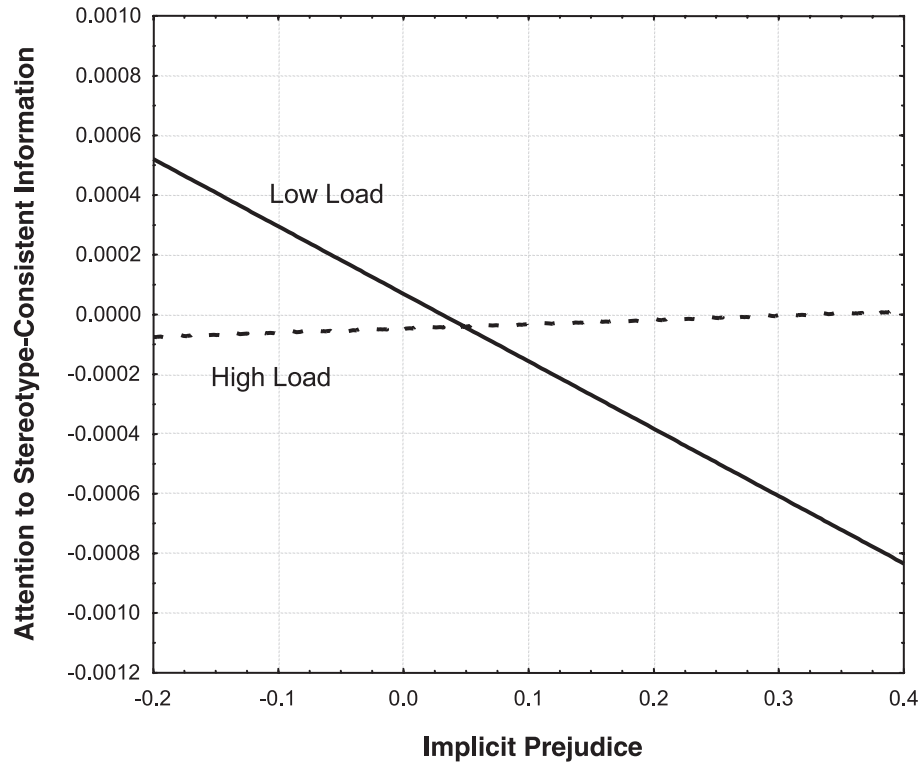


Figure 4. Regression lines predicting attention to stereotype-consistent versus stereotype-inconsistent information as a function of implicit prejudice and cognitive load in Experiment 3. Higher scores on the y-axis indicate greater attention to consistent information. Higher scores on the x-axis indicate a relatively stronger association between Black men and negativity.

iment 3) measure of prejudice. Results from Experiment 2 suggest that these effects may be related to the types of attributions high-prejudice persons make for consistent and inconsistent behaviors. In particular, more external attributions were made for inconsistent than consistent behaviors, and more internal attributions were made for consistent than inconsistent behaviors. These data indicate that high-prejudice people set a higher threshold for accepting counterstereotypic than stereotypic behavior and scrutinize the former much more carefully than the latter (Ditto et al., 2003; Ditto & Lopez, 1992) when they possess the resources to do so. Finally, the correlational data from Experiment 1 and the trait ratings from Experiment 2 indicate that high-prejudice participants were not basing their judgments on the targets' behaviors. Rather, their judgments appear to be driven largely by their preexisting stereotypes. Thus, higher prejudice was associated with both biased encoding and judgment processes, though the latter did not depend on the former.

As for those with low levels of prejudice, there was little evidence that they encoded information in a stereotype-disconfirming fashion. In both Experiments 1 and 2, low-prejudice participants demonstrated no differential processing of consistent and inconsistent information. Only in Experiment 3 was there any evidence that lower levels of prejudice were associated with encoding consistent information more thoroughly than inconsistent information. However, this effect held only in the low-load con-

dition. Despite the implication that low prejudice is associated with attending carefully to consistent information to explain it away (at least, when resources are available), the results of Experiment 2 showed that low-prejudice persons did not differentially explain stereotypic and counterstereotypic behavior. The possibility remains that the low-prejudice participants in Experiment 3 simply were more concerned with disconfirming the stereotype of Black men than the low-prejudice participants in Experiments 1 and 2 were with disconfirming the stereotype of gay men. Further research is needed to delineate the conditions under which low-prejudice people do and do not engage in motivated processing of stereotype-relevant information.

The results from Experiments 1 and 2 also demonstrate that low-prejudice participants based their judgments on the target's specific behaviors. In Experiment 1, low-prejudice participants' target judgments were correlated with the information that was encoded most thoroughly. In Experiment 2, low-prejudice participants readily distinguished between targets demonstrating largely stereotypic and largely counterstereotypic behavior. Thus, low-prejudice participants demonstrated neither biased encoding nor judgment processes, and the two processes were related.

Inhibition, Encoding, and Integration Processes

One important conclusion from these findings is that consideration of the influence of prejudice on the cognitive mechanisms of

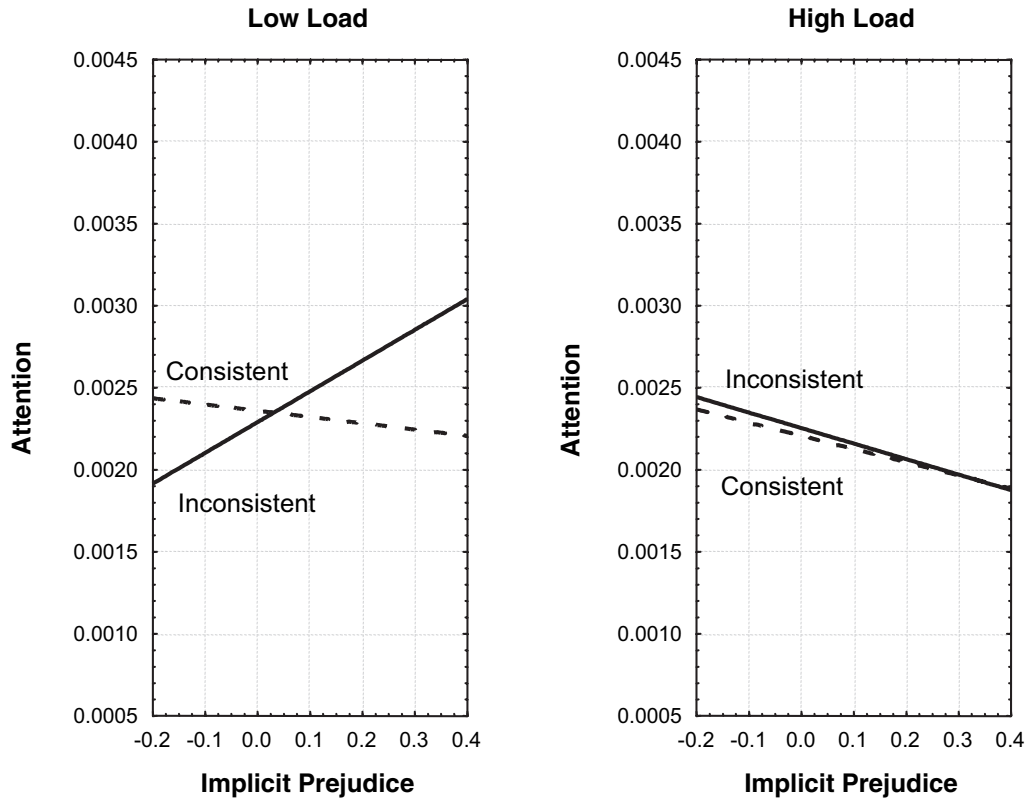


Figure 5. Regression lines predicting attention to stereotype-consistent and stereotype-inconsistent information as a function of implicit prejudice in the low-load (left panel) and high-load (right panel) conditions in Experiment 3. Higher scores indicate greater attention.

stereotype maintenance can extend beyond stereotype-inhibition processes. Prior research has well documented that low-prejudice people are more willing and/or able to inhibit their stereotypes than are high-prejudice people (e.g., Devine, 1989; Monteith, 1993; Monteith et al., 1998). In the present studies, prejudice also has been shown to influence processes as diverse as attention, interpretation, and integration of behavioral information. These results suggest that low-prejudice people also are more willing and/or able to process information about members of the target group in a less biased and more systematic fashion. It does not appear to be the case, however, that low-prejudice people consistently engage in counterstereotypical processing biases. Low-prejudice participants in our studies largely encoded consistent and inconsistent information in the same fashion. What low-prejudice participants did was process the behaviors in an evenhanded, unbiased way. Moreover, the results from Experiments 1 and 2 show that low-prejudice but not high-prejudice people form behavior-based judgments. Thus, they process the behaviors in an unbiased fashion and then use those behaviors to inform their impressions. As such, even in the absence of stereotype inhibition, by processing and integrating a target's behaviors in a nonbiased fashion, low-prejudice people may avoid forming stereotyped impressions.

This analysis also raises a possibility not addressed by inhibition research. Specifically, if a target reliably performs stereotype-consistent behaviors, low-prejudice people may well form stereo-

typical impressions. We would argue, however, that such impressions, while being consistent with the stereotype, do not reflect stereotyping, per se, because they are based on the target's behaviors and not on some prior expectancy. In contrast, impressions of high-prejudice people are likely to be stereotypical regardless of a target's behavior. By failing to consider the implications of a target's behaviors, high-prejudice persons all but cede their impressions to prior expectancies.

We do not wish to suggest that inhibition processes are mutually exclusive with encoding and integration processes. To the contrary, an ability to successfully inhibit a stereotype will likely facilitate efforts at nonbiased and systematic encoding and integration. At the same time, nonbiased and systematic encoding and integration is likely to facilitate inhibition. Consequently, it is difficult to tell how large a role inhibition processes played in producing the present results. Certainly, the data from Experiment 2 indicate that our results were not due to simple differences in the extent to which low and high-prejudice participants inhibited their stereotypes, in that the two groups formed equally stereotypical impressions of the stereotypical target. Thus, when the target was relatively stereotypical, low-prejudice participants did not hesitate to say so. Of course, it is still possible that low-prejudice but not high-prejudice participants inhibited their stereotypes when confronted with the counterstereotypical target. Investigations of the

interactions between inhibition and encoding and integration processes should prove to be an important domain of future research.

Prejudice, Processing Motives, and Attention

Another important goal for future research will be to further specify the relationships among prejudice, processing motives, and attention. The recognition results from Experiment 1 are consistent with those reported elsewhere (Bardach & Park, 1996; Förster et al., 2000; Wänke & Wyer, 1996), showing that participants least motivated to form accurate impressions were most likely to remember counterstereotypical behaviors. These findings (and those from Experiment 3) are consistent with the view that people attend very carefully to information that challenges desired beliefs but give little scrutiny to preference-consistent information (e.g., Ditto & Lopez, 1992; Ditto et al., 2003; Eagly et al., 2000). In contrast, these results do not support the notion that increases in accuracy motivation are associated with attentional biases favoring inconsistent information and that those individuals lacking in accuracy motivation will filter out and ignore information that challenges their beliefs (at least if processing capacity is sufficient). However, there are at least two important differences between our experiments and the experiments that have demonstrated such a relationship between accuracy motivation and attention that may account for the different results (e.g., Erber & Fiske, 1984; Fiske et al., 1999; Fiske & Neuberg, 1990). First, in prior research on the influence of accuracy motivation, motivation has been manipulated primarily through instructions or roles given to the participants. In the absence of these manipulations, participants were not particularly motivated to either individuate or uphold their expectancies; they were relatively neutral observers. In part, this is because the expectancies used in this research have typically consisted of experimenter-induced trait expectancies or stereotypes about groups that fail to evoke strong prejudiced responses (e.g., individuals with schizophrenia; but see Goodwin, Gubin, Fiske, & Yzerbyt, 2000). In contrast, beliefs about homosexuals and Black men are much more volatile and likely to initiate motivated processing, even in the absence of particular instructions. Of particular relevance, the “unmotivated” (high-prejudice) participants in our research may have been motivated to defend their stereotypes, thereby enhancing the desirability of attending to and attempting to discount inconsistent behaviors (as reflected in the attributional data from Experiment 2).

Second, all of the participants in the present experiments were asked to form impressions of the target. In contrast, the presence or absence of such instructions has commonly been manipulated in past research to influence the extent to which participants are motivated to form accurate impressions. Therefore, in our experiments, all participants were likely somewhat motivated by the instructions to attend carefully to the presented information. Perhaps in the absence of such instructions, high-prejudice participants would have been more content to simply ignore the inconsistent behaviors, despite being motivated to defend their stereotypes. Likewise, perhaps low-prejudice participants would have been more likely to show a bias toward disconfirmation and ignore the consistent behaviors. Thus, there are important differences between our experiments and past research that may account for the different relationship we observed between accuracy mo-

tivation (as reflected in levels of prejudice) and the encoding of consistent and inconsistent information than has been previously reported.

Finally, it simply may be the case that there are different varieties of accuracy and defense motivation that influence perceivers in different ways. Thus, the push toward accuracy that derives from outcome dependency or accountability may engender different processes than the accuracy motivation that derives from a personal motivation to avoid seeing others in stereotypical ways (i.e., low prejudice). Likewise, the desire to uphold stereotypes may be played out in different ways, depending on whether that desire derives from a position of power or from a personal animosity toward members of a stereotyped group.

On the Relationships Among Different Measures of Stereotyping and Individuation

More generally, our research raises some important issues about the relationships among different measures of stereotyping and individuation. In the stereotyping literature, encoding, integration, and judgment data all have been used to characterize the extent to which impressions are relatively category based or individuated. For example, a tendency to attend to consistent rather than inconsistent information has been taken as evidence for a category-based rather than a behavior-based impression formation process (e.g., Brewer, 1988; Fiske & Neuberg, 1990). By the same token, a tendency to attend to inconsistent rather than consistent information has been taken as evidence for behavior-based rather than stereotype-based impression formation. It also is sometimes assumed that the content of perceivers' judgments is directly related to the content of the information to which they attend (e.g., Fiske & Neuberg, 1990). However, our data make it clear that the relationships among stereotyping and the different encoding and integration processes are more complicated than has sometimes been acknowledged.

First, attention is a poor predictor of integration versus category use. In Experiments 1 and 3, it was the high-prejudice rather than the low-prejudice participants who were most likely to preferentially attend to inconsistent behaviors. At the same time, the low-prejudice but not the high-prejudice participants formed impressions on the basis of behavioral integration, as reflected in the judgment and correlational data in Experiment 1 and the judgment data in Experiment 2. This shows that attentional data should not be viewed as synonymous with direct evidence of behavioral integration and that individuation does not necessarily involve the preferential encoding of inconsistent over consistent information.

The relationship between attention and judgment stereotypicality appears to be similarly complex. Our results demonstrate a clear dissociation between attention and judgment. Though the high-prejudice but not low-prejudice participants preferentially encoded inconsistent information, the high-prejudice participants made more stereotypical judgments. A similar dissociation was reported by Sherman et al. (1998), who showed that depleting participants' processing resources enhanced the relative attention paid to inconsistent versus consistent information and enhanced the likelihood of drawing stereotypical inferences at the same time. Thus, in considering the relationship between attention and judgment, it is necessary to know not only what information perceivers

are attending to but also why they are attending to that information and how they are encoding it. If high-prejudice perceivers attend carefully to inconsistent information so that they may explain it away, then a simple positive relationship between attention and judgment should not be expected.

Finally, our findings suggest that the relationship between behavioral integration processes and judgmental outcome is not as straightforward as it has been portrayed to be. It is sometimes assumed that the process of integrating a target's behavioral data into an impression will necessarily produce a nonstereotypical judgment. However, we have shown that integration can produce judgments that appear to endorse stereotypes, even among low-prejudice participants. When behavior, in fact, is consistent with stereotypes, then even integration processes among low-prejudice people should produce a stereotype-consistent judgment.

Conclusion

The nature of prejudice has been a consistent focus of social psychological theorizing. Despite this long interest in prejudice, relatively little has been known regarding the specific encoding processes involved in the development and maintenance of negative evaluations of groups and their members. Our studies provide some preliminary evidence of the ways in which individuals high in prejudice engage in different attentional, attributional, and judgment processes than those low in prejudice. However, the relationships between these processes appear to be complex, and a simple characterization of the effects of prejudice remains elusive. The fact that prejudice involves multiple processes occurring together to maintain negative attitudes might help account for the typical resistance of prejudice to challenge or change.

References

- Allport, G. W. (1954). *The nature of prejudice*. Oxford, England: Addison Wesley.
- Bardach, L., & Park, B. (1996). The Effect of in-group/out-group status on memory for consistent and inconsistent behavior of an individual. *Personality and Social Psychology Bulletin*, 22, 169–178.
- Bradley, B. P., Mogg, K., White, J., Groom, C., & de Bono, J. (1999). Attentional bias for emotional faces in generalized anxiety disorder. *British Journal of Clinical Psychology*, 38, 267–278.
- Brewer, M. B. (1988). A dual process model of impression formation. *Advances in Social Cognition*, 1, 1–36.
- Crandall, C. S., & Eshleman, A. (2003). A justification suppression model of the expression and experience of prejudice. *Psychological Bulletin*, 129, 414–446.
- 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.
- Ditto, P. H., & Lopez, D. F. (1992). Motivated skepticism: Use of differential decision criteria for preferred and nonpreferred conclusions. *Journal of Personality and Social Psychology*, 63, 568–584.
- Ditto, P. H., Munro, G. D., Apanovitch, A. M., Scepansky, J. A., & Lockhart, L. K. (2003). Spontaneous skepticism: The interplay of motivation and expectation in responses to favorable and unfavorable medical diagnoses. *Personality and Social Psychology Bulletin*, 29, 1120–1132.
- Eagly, A. H., & Chaiken, S. (1993). *The psychology of attitudes*. Orlando, FL: Harcourt, Brace.
- Eagly, A. H., Kulesa, P., Brannon, L. A., Shaw, K., & Hutson-Comeaux, S. (2000). Why counterattitudinal messages are as memorable as proattitudinal messages: The importance of active defense against attack. *Personality and Social Psychology Bulletin*, 26, 1392–1408.
- Erber, R., & Fiske, S. T. (1984). Outcome dependency and attention to inconsistent information. *Journal of Personality and Social Psychology*, 47, 709–726.
- Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*, 69, 1013–1027.
- Fiske, S. T., Lin, M., & Neuberg, S. L. (1999). In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 231–254). New York: Guilford Press.
- Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuated processes: Influences of information and motivation on attention and interpretation. *Advances in Experimental Social Psychology*, 23, 1–74.
- Förster, J., Higgins, E. T., & Strack, F. (2000). When stereotype disconfirmation is a personal threat: How prejudice and prevention focus moderate incongruity effects. *Social Cognition*, 18, 178–197.
- Frey, D. (1986). Recent research on selective exposure to information. *Advances in Experimental Social Psychology*, 19, 41–80.
- Gilbert, D. T., Pelham, B. W., & Krull, D. S. (1988). On cognitive busyness: When person perceivers meet persons perceived. *Journal of Personality and Social Psychology*, 54, 733–740.
- Goodwin, S. A., Gubin, A., Fiske, S. T., & Yzerbyt, V. Y. (2000). Power can bias impression processes: Stereotyping subordinates by default and by design. *Group Processes and Intergroup Relations*, 3, 227–256.
- Graesser, A. C. (1981). *Prose comprehension beyond the word*. New York: Springer-Verlag.
- Grier, J. B. (1971). Nonparametric indexes for sensitivity and bias: Computing formulas. *Psychological Bulletin*, 75, 424–429.
- Haddock, G., Zanna, M. P., & Esses, V. M. (1993). Assessing the structure of prejudicial attitudes: The case of attitudes toward homosexuals. *Journal of Personality and Social Psychology*, 65, 1105–1118.
- Hamilton, D. L., & Sherman, J. W. (1994). Stereotypes. In R. S. Wyer Jr., & T. K. Srull (Eds.), *Handbook of social cognition* (2nd ed., Vol. 2, pp. 1–68). Mahwah, NJ: Erlbaum.
- Hastie, R. (1984). Causes and effects of causal attribution. *Journal of Personality and Social Psychology*, 46, 44–56.
- Herek, G. M. (Ed.). (1998). *Stigma and sexual orientation: Understanding prejudice against lesbians, gay men, and bisexuals*. Thousand Oaks, CA: Sage.
- Jacoby, L. L., Toth, J. P., & Yonelinas, A. P. (1993). Separating conscious and unconscious influences of memory: Measuring recollection. *Journal of Experimental Psychology: General*, 122, 139–154.
- Jost, J. T., & Hunyady, O. (2002). The psychology of system justification and the palliative function of ideology. *European Review of Social Psychology*, 13, 111–153.
- Jost, J. T., & Major, B. (Eds.). (2001). *The psychology of legitimacy: Emerging perspectives on ideology, justice, and intergroup relations*. New York: Cambridge University Press.
- Katz, D., & Braly, K. (1933). Racial stereotypes of one hundred college students. *Journal of Abnormal & Social Psychology*, 28, 280–290.
- Kay, A. C., & Jost, J. T. (2003). Complementary justice: Effects of “poor but happy” and “poor but honest” stereotype exemplars on system justification and implicit activation of the justice motive. *Journal of Personality and Social Psychology*, 85, 823–837.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108, 480–498.

- Larsen, K. S., Reed, M., & Hoffman, S. (1980). Attitudes of heterosexuals toward homosexuality: A Likert-type scale and construct validity. *Journal of Sex Research, 16*, 245–257.
- Lippmann, W. (1922). *Public opinion*. New York: Harcourt, Brace.
- Mackie, D. M., & Smith, E. R. (1998). Intergroup relations: Insight from a theoretically integrative approach. *Psychological Review, 105*, 499–529.
- 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., Sherman, J. W., & Devine, P. G. (1998). Suppression as a stereotype control strategy. *Personality and Social Psychology Review, 2*, 63–82.
- Nosek, B. (2004). *Moderators of the relationship between implicit and explicit attitudes*. Unpublished Manuscript.
- Nosek, B., & Banaji, M. R. (2001). The Go/No-go Association Task. *Social Cognition, 19*, 625–666.
- Petty, R. E., & Wegener, D. T. (1998). Attitude change: Multiple roles for persuasion variables. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *Handbook of social psychology* (4th ed., Vol. 1, pp. 323–390). New York: McGraw-Hill.
- Plaks, J. E., Stroessner, S. J., Dweck, C. S., & Sherman, J. W. (2001). Person theories and attention allocation: Preferences for stereotypic versus counterstereotypic information. *Journal of Personality and Social Psychology, 80*, 876–893.
- Plant, E. A., & Devine, P. G. (1998). Internal and external motivation to respond without prejudice. *Journal of Personality and Social Psychology, 75*, 811–832.
- Richardson-Klavehn, A., & Bjork, R. A. (1988). Measures of memory. *Annual Review of Psychology, 39*, 475–543.
- Sherman, J. W., Conrey, F. R., & Groom, C. J. (2004). Encoding flexibility revisited: Evidence for enhanced encoding of stereotype-inconsistent information under cognitive load. *Social Cognition, 22*, 214–232.
- Sherman, J. W., & Frost, L. A. (2000). On the encoding of stereotype-relevant information under cognitive load. *Personality and Social Psychology Bulletin, 26*, 26–34.
- Sherman, J. W., Lee, A. Y., Bessenoff, G. R., & Frost, L. A. (1998). Stereotype efficiency reconsidered: Encoding flexibility under cognitive load. *Journal of Personality and Social Psychology, 75*, 589–606.
- Strull, T. K. (1981). Person memory: Some tests of associative storage and retrieval models. *Journal of Experimental Psychology: Human Learning and Memory, 7*, 440–462.
- Strull, T. K. (1984). Methodological techniques for the study of person memory and social cognition. In R. S. Wyer Jr. & T. K. Strull (Eds.), *Handbook of social cognition* (1st ed., Vol. 2, pp. 1–72). Hillsdale, NJ: Erlbaum.
- Strull, T. K., Lichtenstein, M., & Rothbart, M. (1985). Associative storage and retrieval processes in person memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11*, 316–345.
- von Hippel, W., Sekaquaptewa, D., & Vargas, P. (1997). The linguistic intergroup bias as an implicit indicator of prejudice. *Journal of Experimental Social Psychology, 33*, 490–509.
- Wänke, M., & Wyer, R. S., Jr. (1996). Individual differences in person memory: The role of sociopolitical ideology and in-group versus out-group membership in responses to socially relevant behavior. *Personality and Social Psychology Bulletin, 22*, 742–754.
- Yzerbyt, V., & Anouk, R. (2001). Blame it on the group: Entitativity, subjective essentialism, and social attribution. In J. T. Jost & B. Major (Eds.), *The psychology of legitimacy: Emerging perspectives on ideology, justice, and intergroup relations* (pp. 103–134). New York: Cambridge University Press.
- Yzerbyt, V., Rocher, S., & Schadrin, G. (1997). Stereotypes as explanations: A subjective essentialistic view of group perception. In R. Spears & P. J. Oakes (Eds.), *The social psychology of stereotyping and group life* (pp. 20–50). Malden, MA: Blackwell Publishers.

Received July 17, 2003

Revision received April 14, 2005

Accepted April 14, 2005 ■