Journal of Personality and Social Psychology
2005, Vol. 89, No. 4, 488–503

0022-3514/05/$12.00 DOI: 10.1037/0022-3514.89.4.488

Conceptualizing Control in Social Cognition: How Executive Functioning Modulates the Expression of Automatic Stereotyping

B. Keith Payne
Ohio State University

Two studies investigated the role of executive control in moderating the relationship between automatic stereotype activation and behavioral responses. Race bias in weapon identification was used to measure stereotyping, and a process dissociation procedure was used to measure automatic and controlled components of performance. In Experiment 1, the controlled component was shown to correlate with general attention control and race-specific motivations to control prejudice. Across multiple measures, automatic race bias was more likely to be expressed as behavioral discrimination among individuals with poor executive control. Experiment 2 found the same relationship between automatic and controlled components of behavior when predicting impressions of a Black individual. Executive control is discussed in the context of other control strategies in influential dual-process models of stereotyping.

Keywords: implicit, attitudes, automatic, projective

The modern history of racial prejudice in America is largely a history of conflict between egalitarian ideals and ingrained inequalities. From Myrdal's (1944) consideration of the "American dilemma" and Allport's (1954) description of "prejudice with compunction" to contemporary dual-process theories of prejudice, the disconnection between intentions and behavior has occupied center stage. Theories aimed at explaining the psychology of prejudice are therefore increasingly emphasizing the distinction between intentionally controlled and unintentional aspects of behavior. This interplay between automatic and controlled aspects of race bias is the topic of the present research.

This article focuses on the role of executive control in determining whether an automatically activated bias becomes expressed in overt behavior. Executive control is the capacity to constrain thought processes and behavior to reach goal-relevant ends. Executive control refers to a constellation of interrelated mental processes that bring about this capacity, rather than a single unitary process. These processes include planning and monitoring behavior; coordinating behavior in complex, novel, or ambiguous situations; selectively activating information that facilitates one's goal, while actively inhibiting information that interferes; and overriding impulsive or automatic responses when they clash with goals (see Baddeley, 1986; Norman & Shallice, 1986).

Social cognition research has emphasized the role of limited-capacity cognitive resources in a wide range of topics (for examples, see Chaiken & Trope, 1998). However, advances in cognitive science and cognitive neuroscience have provided a more fine-grained analysis of the operations involved in cognitive control processes, beyond the general idea of "resources" (Logan, 1997; Baddeley, 1986). For example, an influential model of working memory distinguishes between domain-specific "slave systems," which operate specifically on verbal or visual–spatial material, and a central executive system, which regulates them. Each of these systems has its own domain-specific capacity limitations, but they do not necessarily interfere with each other because they do not draw on the same generic "resource pool" (Baddeley, 1986). By more closely specifying the cognitive operations underlying control processes, this article attempts to shed more light on the ways that the mind regulates the impacts of automatic biases. The present research focuses on aspects of executive control related to the selective allocation of attention and the inhibition of impulsive responses.

Social psychological research has recently begun to address the ways in which executive control processes influence social judgments and social behavior (Bodenhausen & Macrae, 1998; Macrae, Bodenhausen, Schloerscheidt, & Milne, 1999). For example, von Hippel, Silver, and Lynch (2000) showed that differences in inhibitory ability—a critical component of executive control—mediated the tendency for older adults to express greater racial prejudice than younger adults (see also von Hippel & Gonsalkorale, 2005).

Research by Richeson and Shelton (2003) has also illustrated the importance of executive functioning in social regulation. After an interaction with a Black partner, White research participants showed impaired performance on a measure of executive control. This impairment was greatest for those participants who earlier showed strong racial bias on an implicit measure of race attitudes. These results were interpreted as reflecting a depletion of execu-

B. Keith Payne, Department of Psychology, Ohio State University.

This research was conducted as part of the author's doctoral dissertation under the supervision of Alan Lambert and Larry Jacoby. I would like to thank committee members Michael Strube, David Balota, Pascal Boyer, and Keith Sawyer. Portions of this research were funded by National Institute of Mental Health Grant 1-R03-MH070573-01 awarded to B. Keith Payne. Thanks finally to Russ Fazio, Elizabeth Marsh, and Brandon Stewart for comments and discussions of earlier drafts of the article.

Correspondence concerning this article should be addressed to B. Keith Payne, who is now at the Department of Psychology, University of North Carolina, Davie Hall, CB #3270, Chapel Hill, NC 27599. E-mail: payne@unc.edu

tive control capacity because of the self-regulatory demands of interacting with a Black partner.

Efforts to regulate automatic social biases have been linked with brain regions critical to executive functioning. For instance, Cunningham et al. (2004) found that when Black and White faces were flashed too quickly to be consciously detected, Black faces evoked stronger amygdala responses than White faces. However, when the faces were presented more slowly, this difference in amygdala response was reduced, and Black faces evoked stronger activation in areas of the prefrontal cortex and the anterior cingulate cortex associated with executive functions (see also Richeson et al., 2003, and Amodio et al., 2004, discussed below).

Together, these lines of research highlight the importance of basic executive control processes in regulating prejudiced behavior. The present work develops and extends this theme by proposing a conceptual framework along with an empirical measurement model for relating three important aspects of prejudice: automatically activated attitudes and stereotypes, basic executive control processes, and discriminatory behavior. The broad thesis advanced here is that behavioral expression of automatically activated stereotypes depends on how strongly executive control is engaged. As will be discussed in more detail, differences in executive control may arise because of differences in ability or because of differences in motivation and effort. Both of these factors may vary chronically or from one situation to the next. Because of the critical role of executive control in overriding automatic or impulsive responses, planning and executing complex and subtle behaviors, and coping with novelty and ambiguity, this class of cognitive process is expected to be a critical factor in determining when automatic stereotyping translates into actual discrimination.

Process Dissociation and Dual-Process Theories of Social Cognition

The research strategy for the current experiments uses the process dissociation procedure (PDP; Jacoby, 1991) as a unifying framework. This section describes that procedure and places it in the theoretical context of dual-process theories in social psychology. Recent dual-process theories of prejudice have identified conditions under which behavior is likely to be driven by intentionally controlled versus automatic processes (e.g., Devine, 1989; Fazio, 1990). For example, Devine's (1989) dissociation model specifically addresses the disconnection between automatically activated prejudices and intended responses. Individuals motivated to be unprejudiced often engage in controlled processes in an attempt to prevent prejudiced actions. However, even individuals with egalitarian values may sometimes find prejudice seeping into their behaviors (Devine, Monteith, Zuwerink, & Elliot, 1991). Under conditions in which people are not monitoring their behaviors or in which intentional control is difficult such as distraction, fatigue, or fast responding, unintended prejudice may cause acts of discrimination (Bodenhausen, 1990; Macrae, Milne, & Bodenhausen, 1994).

One popular method used to separate the contributions of automatic and controlled processes in race biases within dual-process frameworks is to compare performance on implicit and explicit tasks. Commonly used implicit tasks include implicit association tests (IATs; Greenwald, McGhee, & Schwartz, 1998), priming tasks (Fazio, Jackson, Dunton, & Williams, 1995; Wittenbrink, Judd, & Park, 1997) and word completions (Gilbert & Hixon, 1991). These tasks allow researchers to infer the activation of stereotypes or attitudes from changes in participants' performance on an objective task, rather than relying on self-reports. Explicit tasks such as questionnaires and trait rating methods directly ask participants to report on their attitudes and stereotypes or to make a stereotype-relevant judgment. Comparisons between implicit and explicit tasks have proven very useful because performance on the two classes of measures has frequently been dissociated, suggesting that implicit cognitive processes and strategically expressed reports may follow different principles. For example, Fazio and colleagues (1995) found that the presence of a Black experimenter elicited more positive racial attitudes on a self-report scale but did not alter performance on a priming measure of racial attitudes.

Nonetheless, one problem is that comparing implicit and explicit tasks leaves the processes of interest confounded with the properties of the measurement instrument. For example, self-report measures could be dissociated from reaction time measures for reasons related to the methods rather than the process of interest. A dissociation could potentially be due to differences in the sensitivities of the tests, different response scales (e.g., rating scales vs. response latencies), different reliabilities for the two tasks, and so on. A second limitation is that no task is likely to be "process pure," because various admixtures of automaticity and control may feed into any given behavior (e.g., Bargh, 1989; see Payne, Jacoby, & Lambert, 2005 for a discussion of this and other strengths and limitations of the task dissociation approach).

In contrast to task dissociation methods, the process dissociation procedure (PDP) was developed to separate automatic and controlled contributions to the same task (Jacoby, 1991). The PDP framework begins with the assumption that both automatic and controlled processes contribute to any given behavior. The relative contribution of each process can be estimated by arranging tasks in which automatic and controlled processes are placed both in opposition and in concert. For example, our research group has used process dissociation to separate automatic and controlled processing in a weapon identification task. In this task, a photo of either a Black or White face precedes a photo of a gun or tool; the task is to ignore the face and press one key for "gun" and another for "tool" (Payne, 2001). When a Black face precedes a gun, unintended race biases and the intended "gun" response both lead to the same behavior (a congruent trial). In contrast, when a Black face precedes a tool, the unintended race bias favors the "gun" response, whereas carrying out the intended identification process favors the "tool" response (an incongruent trial). When both sources of information favor the same response, then correct responses could result from either controlled responding (C) or automatic bias (A) given the failure of control (1 − C). This relationship may be expressed mathematically in the equation:

$$P(\text{correct}|\text{congruent}) = \text{C} + \text{A}(1 - \text{C}) \qquad (1)$$

That is, when a Black face was paired with a gun, the correct response ("gun") could be consistently achieved in two ways. The first way was by successfully controlling the response. The second was by unintentionally responding "gun" because of the race prime, even when unable to implement the response based on the actual target.

On incongruent trials, the two processes were set in opposition to one another. When cognitive control and stereotypic automatic bias were opposed to one another, false alarms would occur when an automatic bias (A) operated in the absence of control (1 – C). Mathematically, this can be written as:

$$P(\text{stereotypic}|\text{incongruent}) = A(1 - C) \qquad (2)$$

As an example, consider the trials in which a Black face was paired with a tool. If control failed and a participant was influenced by activated stereotypes, then he or she would respond "gun." Given these equations, one can solve algebraically for estimates of cognitive control (C) and automatic bias (A):

$$C = P(\text{correct}|\text{congruent}) - P(\text{stereotypic error}|\text{incongruent}) \qquad (3)$$

$$A = P(\text{stereotypic error}|\text{incongruent})/(1 - C) \qquad (4)$$

It should be noted that these algebraic estimates involve a number of assumptions, some of which have been controversial. In the General Discussion section, these assumptions and their controversies will be considered in light of the results of the present research.

When the PDP estimates were obtained for performance in the weapon misidentification task, the data revealed striking dissociations. For example, requiring participants to respond quickly increased the rate of stereotypical errors (Payne, 2001). This increase was driven by a reduction in the controlled component under speeded conditions, with no effect on the automatic component (see also Lambert et al., 2003). Other research has shown that making race salient can increase stereotypic misidentifications by increasing the automatic component, without affecting the controlled component (Payne, Lambert, & Jacoby, 2002). Taken together, these results begin to create a coherent narrative about how stereotypical actions can be driven by either changes in the automatic or the controlled components of performance.

The story is incomplete, however, without a full understanding of the psychological processes underlying the PDP estimates. The PDP control estimate is one type of discriminability parameter (similar to d' in signal detection theory, or SDT; Tanner & Swets, 1954). Note that in itself, a discriminability parameter does not reveal which process is responsible for the discriminations. Discriminability estimates can be calculated for virtually any kind of measure for which there are correct and incorrect answers. The psychological meaning of the estimate depends entirely on the task. For example, discriminability estimates may be calculated on intelligence tests, eye examinations, or calling bluffs during poker games. Each, of course, reflects very different psychological processes.

What does the PDP control estimate represent in stereotyping tasks such as the weapon identification task? In this task, it literally quantifies participants' ability to behaviorally distinguish between weapons and tools, correcting for response biases—but how to label the underlying psychological process is controversial. Simply labeling the process as "control" is not satisfactory by itself. For example, various dual-process models include a number of different strategies under the rubric "control," including how people may gather individuating information, replace biased cognitive and behavioral responses with more egalitarian ones, correct their judgments by adjusting their response according to their lay theories of bias, or attempt to suppress stereotype-relevant thoughts (Devine & Monteith, 1999).

This article provides evidence for a straightforward explanation of the control estimate: that it represents executive control processes. The description of control outlined here shares some features with other varieties of controlled processes frequently described in social psychology. For example, constraining one's processing to goal-relevant information bears some similarity to basing judgments on individuating information. However, the two are clearly not identical. Responding as intended while avoiding interference bears some similarity to ideas of suppression or inhibition. However, cognitive control cannot fully be described by what is kept out of mind; what is selected and maintained in mind is also important. Although the PDP definition of control differs from the control strategies usually discussed in social psychology, it has a great deal of precedent in cognitive and cognitive–neuroscience literatures on selective attention and executive function (Baddeley, 1986; Hasher & Zacks, 1988; Kane, Bleckley, Conway, & Engle, 2001).

The most direct evidence currently available that the PDP estimate of control in the weapon misidentification task reflects executive control comes from recent work using event-related potentials (ERPs; Amodio et al., 2004). This study found that an ERP signal known as error-related negativity (ERN) was associated with response control in the weapons task. The ERN is a signal believed to be generated in the anterior cingulate cortex, an area of the frontal cortex heavily implicated in the strategic planning and control of behavior (Botvinick, Nystrom, Fissel, Carter, & Cohen, 1999). The ERN is typically found when people process information that creates conflicts that need to be resolved, as in cases of response competition.

According to a neuropsychological model of executive control, conflict detection is an important prerequisite for implementing executive control (Botvinick, Braver, Barch, Carter, & Cohen, 2001; Botvinick et al., 1999). As predicted by this model, the ERN signal correlated significantly with PDP estimates of controlled, but not automatic, components. Other studies have shown the ERN to be similarly correlated with performance on selective attention tasks thought to reflect executive control processes, such as Stroop and flanker tasks (Botvinick et al., 1999). This evidence suggests that response control in selective attention tasks and weapon identification tasks may rely on similar neural processes.

Each of the control strategies reviewed above—individuation, correction, and suppression—cannot be reduced to executive control, because each is a complex ensemble of processes. Yet each involves at a basic level the ability to keep one planned set of processes "on track" and to block out another set. They all involve the need to separate relevant inputs from irrelevant inputs and to maintain focus on the goal that defines some information as more relevant than other information. Executive control might be a fundamental building block that makes these various control strategies possible (see also Macrae et al., 1999). Why should those interested in automatic social cognition and behavior be interested in executive control? Because executive control describes the basic attentional capacities that may determine whether or not automatic processes drive behavior.

Overview of Primary Hypotheses

Two studies are presented here to test three primary hypotheses. The first hypothesis is that the PDP control estimate, in the context of the weapon identification task described above, provides a valid measure of the tendency to exert control over the effects of automatically activated stereotypes and attitudes. A corollary of this hypothesis is that the control estimate should be related to both general executive control abilities and domain-specific motivations. In one of the seminal works framing attentional control as a limited-capacity resource, Kahneman (1973) argued that although attentional capacity is finite, its boundaries are flexible. Factors such as motivation, effort, and arousal can increase the size of the "pool," at least temporarily. In this light, it was expected that participants who had higher general executive ability would display greater control in a stereotyping task because of general skill or ability levels. In addition, it was expected that participants who were strongly motivated to avoid acting with prejudice would also exert greater control on the stereotyping task. This is not because they would have higher ability in general, but because they would be highly motivated to exert effort to avoid discriminating.

The second hypothesis was that the PDP automatic estimate would be related to other measures of implicit attitudes. This is a straightforward implication of the idea that the automatic estimate in the weapons task captures automatically activated attitudes and stereotypes.[1]

Finally, the third hypothesis was that executive control would moderate the relationship between automatically activated prejudice and overt behavioral discrimination. In the first experiment, this hypothesis was tested using misidentifications in the weapons task as the measure of discriminatory behavior. In the second, the focus was on a much more removed act of impression formation. In the first case, executive control was expected to be critical for overcoming activated but inappropriate responses in the presence of conflicting tendencies (i.e., negative associations to Blacks vs. intentions to respond objectively to weapons and harmless items). In the second case, executive control was expected to be critical for ignoring activated prejudices while coping with a highly ambiguous new person and forming an impression.

Experiment 1: Correlate of Automatic and Controlled Estimates

The first goal of Experiment 1 was to directly test the link between the PDP control estimate and known measures of executive control. The first step was to obtain PDP estimates of cognitive control and automatic bias for each subject. To add generality, these estimates were obtained both from the weapon misidentification procedure and from a second implicit test also capable of yielding the PDP estimates. This second implicit test was a word evaluation task, a modified evaluative priming task containing congruent and incongruent trials.

Only the control estimates (and not the automatic estimates) from the two implicit attitude tests should correlate with performance on the executive control task, an antisaccade test. The antisaccade task measures voluntary attentional control by examining people's ability to control the automatic orienting response. When a new stimulus abruptly appears in a person's visual field, the natural reflexive response is to focus visual attention on it. In the antisaccade task, participants are instructed to resist the powerful orienting reflex by directing their attention away from an abruptly appearing stimulus. A number of neuropsychological studies have shown that patients with prefrontal brain damage are impaired on this task (Everling & Fischer, 1998). Antisaccade task performance is also related to other measures of normal executive functioning such as working memory and Stroop and flanker task performance (Kane et al., 2001). The surface characteristics of this task are very different from those in the weapon identification and word evaluation tasks. However, the argument is that all three tasks require participants to overcome an automatically activated response in order to act in line with their goal. To the extent this is true, control (but not automatic) estimates from the two implicit tests should correlate with performance on the antisaccade task.

Beyond tying the PDP estimate of control to executive processes, Experiment 1 also addresses the question of whether cognitive control actually reduces the impact of automatic biases on performance. The logical basis for this prediction is found in the equations for calculating the process estimates. If the reader looks back at Equation 2, he or she will be reminded that stereotypical errors are not a pure reflection of automatic activation. Instead, stereotypical errors are a joint function of both automatic stereotype activation and failures of cognitive control. This implies that as cognitive control becomes lower, a given level of automatic bias produces a higher number of stereotypical errors (relative to counterstereotypical errors). Therefore, one would expect a strong correlation between the degree of automatic bias and the proportion of stereotypical errors when cognitive control is low. Conversely, one would expect a weaker correlation between the degree of automatic bias and the number of stereotypical errors when cognitive control is high.

There is a parallel here to the ways that controlled processing is often understood in social cognition research. Under conditions in which there is a high degree of controlled processing, automatic processing is not expected to drive behavior. However, the reasons are quite different in the present model compared to "correction" or "deliberative thinking" models of control. In the model outlined here, when control is high, people constrain their processing to the relevant information rather than being driven by irrelevant but activated information. A stereotype, for example, may be highly activated, but that only matters for behavior when cognitive control fails.

---

[1] In the present work, I did not distinguish between evaluative and descriptive associations, because in everyday contexts, the descriptive aspects of racial stereotypes also tend to be unfavorable. The measures used were based primarily on evaluations because there is evidence that prejudice may be a more fundamental predictor of discriminatory behavior than stereotypes (Dovidio et al., 1996). Furthermore, existing research using the weapon identification paradigm provides evidence that both descriptive (i.e., stereotype) and evaluative (i.e., attitude) aspects of group representations figure into the weapon bias (Correll et al., 2002; Judd, Blair, & Chapleau, 2004; Payne, 2001). Still, descriptive and evaluative components are conceptually and empirically distinguishable, and to the extent that they function differently, this is an interesting avenue for future research.

In summary, Experiment 1 tested for specific correlates of automatic and controlled estimates obtained in two implicit attitude measures, the weapon identification task and the word evaluation task. Participants also completed the antisaccade task, a measure of attentional control predicted to be related to the control estimates from the implicit tests, but not to the automatic estimates. Participants also took the IAT (Greenwald et al., 1998), which as an implicit test should correlate with the automatic but not the control estimate (a common factor of automatic racial bias should underlie all three race-based measures). Finally, racial motivations were measured using the Motivation to Control Prejudiced Responses Scale (MCPR; Dunton & Fazio, 1997). This measure should not be related to the antisaccade task (people with more executive control should not necessarily have greater motivation to control prejudice), but it might be related to the control estimates in the two implicit tasks (Amodio et al., 2004). That is, participants most concerned about acting with prejudice may increase their effort on the prejudice-related tasks, thereby increasing the degree of intentional control observed. The degree of response control recruited in the two implicit tasks was expected to be a function of both general ability (as measured by performance on the antisaccade test) and domain-specific motivational factors (as measured by performance of the MCPR).

The measures obtained in Experiment 1 allow not only for the examination of correlations across measures but also an examination of the degree to which cognitive control moderates when automatic biases drive responses. Both scores on the antisaccade task and control estimates from the two implicit tasks were used as moderators, testing whether levels of control determined when automatic stereotyping was expressed as behavioral bias.

Method

Participants

Seventy-six non-Black participants (55 women, 21 men) took part in the experiment for course credit. A computer error caused a loss of data for eight participants on the weapon identification task. As a result, the degrees of freedom reported vary depending on whether this task is included in the analyses.

Measures

The experiment consisted of four computer-based tasks followed by a set of self-report measures and demographic questions. The self-report measure of primary interest was the MCPR (Dunton & Fazio, 1997). The computerized tasks were the weapon identification task, IAT, word evaluation task, and antisaccade task. The IAT and evaluative priming tasks were implemented using response deadlines so that all measures could be computed using the same metric of accuracy rather than comparing accuracy and response time measures. Both the IAT (Cunningham, Preacher, & Banaji, 2001) and the evaluative priming task (Draine & Greenwald, 1998; Musch & Klauer, 2001; Otten & Wentura, 1999) have been studied using response deadline procedures. The results in these studies show that accuracy under response deadline displays the same pattern of facilitation as reaction times: accuracy is higher on compatible trials than on incompatible trials.

At the beginning of the session, the experimenter told participants that they would be asked to complete a number of computer tasks that were described as measures of concentration. The antisaccade task was presented first, both because the task bolstered the cover story that the experiment concerned concentration and because this task was not expected to affect any of the race-related measures to follow. The next three computer tasks were counterbalanced in order. The self-report measures were completed at the end of the session.

Weapon identification task. The weapon identification task was nearly identical to that used in Payne (2001, Study 2). The experimenter described the task as a test of concentration that required both speed and accuracy. Participants were told that they would see pairs of pictures flashed briefly on the monitor. They were instructed to do nothing with the first picture, a face that would serve as a warning signal that the target picture was about to be presented. They were instructed to respond to the second picture, which would always be either a gun or a tool. Participants' task was to classify each target object as either a gun or a tool by pressing one of two keys. The experimenter instructed participants that they should respond to each object very quickly or a red exclamation mark would appear to signal that the response had been too slow.

The first picture (the prime) consisted of a White or Black face. The second picture (the target) was either a handgun or a hand tool. Prime and target stimuli were 5.3×4 cm digital images. The primes included four Black and four White male faces. Target photos included four handguns with varying features and four hand tools (e.g., wrench, pliers). The prime remained on the screen for 200 ms and was replaced immediately by the target. After the target was presented for 200 ms, a visual mask replaced it. The mask remained on the screen until the participant responded or until the feedback mark appeared. If participants responded within the 500-ms deadline, no feedback was given. Before the experimental trials began, participants received a block of 64 practice trials to become acquainted with the task and to practice classifying items quickly and accurately. After the practice trials, participants completed an additional set of 128 critical trials. The computer program randomly ordered the prime–target pairs for each participant.

IAT. The IAT used the same face stimuli used as primes in the weapon identification task. The pleasant and unpleasant words were taken from previous work using the IAT (Greenwald et al., 1998). The positive words were: *fabulous*, *happiness*, *terrific*, and *wonderful*. The negative words were *brutal*, *disgusting*, *nasty*, and *horrible*. As in the task designed by Greenwald and colleagues (1998), the task was divided into eight blocks (each block contained 16 trials). First, participants simply classified Black and White faces, with no word evaluations. Next, they classified words with no face classifications. The third block began the critical trials, with the two classifications combined. In stereotype-consistent blocks, participants were asked to decide whether each stimulus was "Black OR bad" versus "White OR good." In the stereotype-inconsistent blocks, participants were asked whether each stimulus was "Black OR good" versus "White OR bad." As in the procedure described by Greenwald et al. (1998), the blocks alternated between stereotype-consistent and stereotype-inconsistent, with consistent pairings on Blocks 1, 4, 5, and 8. The order of race and words was randomized within each block. A total of eight blocks, including 128 individual trials, was completed. Participants were required to respond within 500 ms. Slow responses triggered a red exclamation point as feedback.

Word evaluation task. The evaluative priming task was based on priming procedures that have been used in many other studies (e.g., Fazio et al., 1995). The primes were the same photos of Black and White men used in the weapon identification task and the IAT, and the target words were the same words used in the IAT. Participants responded by pressing one key for positive words, and another for negative words. One departure from some versions of the task was that primes and targets were presented simultaneously. The face always appeared in the center of the screen. The location of the words alternated randomly above or below the face (within 1 cm). The pair remained on the screen for 500 ms or until the participant made a response. If no response was made within 500 ms, a red exclamation mark signaled that the response was too slow. Each trial was separated from the previous response by a 500-ms interval.

Simultaneous presentation was chosen for two reasons. First, research has shown that evaluative priming effects are largest at extremely short stimulus onset asynchronies (SOA), including simultaneous presentation (zero SOA; Hermans, De Houwer, & Eelen, 2001; Musch & Klauer, 2001). A second reason is that simultaneous presentation highlights the need to ignore one source of information while attending to another. This method makes clearer the parallels between processes involved in evaluative priming and executive control tasks such as the flanker task (Eriksen & Eriksen, 1974).

Participants were instructed that they would see faces and words, and their task was to ignore the faces while evaluating the words. The task was described as an attempt to study how pictorial and verbal information might interfere with each other. Participants were encouraged to respond as quickly and accurately as possible or an exclamation point would signal a slow response. Before the critical trials, participants received 24 practice trials. The critical trials included four blocks, with 64 trials per block, for a total of 256 trials.

Antisaccade task. The antisaccade task measures voluntary attentional control by examining people's ability to control the automatic orienting response. The antisaccade task was designed according to the descriptions in Everling and Fischer (1998) and Kane et al. (2001). The task consisted of two blocks. The first block was a "prosaccade" block, and the second was an "antisaccade" block. That is, on the first block, the target appeared on the same side of the screen as the cue. In the second block, the target appeared on the opposite side. Only performance on the antisaccade block is of interest because performance on the prosaccade block is usually perfect. On each trial, a fixation point first appeared for 2,000 ms. Following the fixation, a red circle served as the cue, appearing in one of two positions on the left or right side of the monitor. Left versus right orientation varied randomly on a trial-by-trial basis. The cue remained on the screen for 400 ms. At that point, the cue disappeared, and the target appeared either on the left or right of the screen in one of the same two positions. The target was always an "H" or a "T." The target remained on the screen for 100 ms and then was masked by the pound symbol (#).

Participants were asked simply to indicate whether the target was an H or a T as quickly as possible. In the antisaccade block, they were instructed that the cue and target would appear on opposite sides of the screen and were instructed to look away from the cue in order to identify the target. Participants completed 48 trials in each of the two blocks. Because of the brief target presentation, if participants looked toward the cue rather than away from it, they ran the risk of not seeing the target on antisaccade trials. Accuracy was used as the dependent variable of interest, and no response deadline was imposed.

MCPR. The MCPR (Dunton & Fazio, 1997) was included to test whether people's motivations to avoid behaving with prejudice would be associated with greater cognitive control on the race-related tasks. This measure includes two subscales. The first is *concern with acting prejudiced*, which reflects a personal internalized desire to not behave in prejudiced ways. The second subscale is *restraint to avoid dispute*, which reflects a more externally driven desire to avoid prejudiced responses for the sake of avoiding conflict.

Procedure

After providing informed consent, participants were seated at a computer terminal. The experimenter described the experiment as a study of concentration and focused attention. As I described earlier, all participants first completed the antisaccade task and then completed the three race-related tasks in a counterbalanced order. After these tasks, participants completed the MCPR, then were fully debriefed, and thanked for their participation.

Results

The hypotheses all concerned individual differences rather than mean performance. Therefore mean results will be summarized briefly and analyses will focus on individual differences. Finally, regression analyses will be reported that tested the interactive relationships between automatic and controlled processes.

The weapon identification task was analyzed using a 2 (Prime Race) $\times$ 2 (Target Object) repeated measures analysis of variance (ANOVA) with error rates as the dependent variable. Replicating previous research with this task, the predicted Prime (race) $\times$ Target (object) interaction was significant, $F(1, 68) = 9.50$, $p < .01$. Participants mistook tools for guns more often after a Black prime ($M = .20$, $SD = .16$) than after a White prime ($M = .15$, $SD = .13$) They also mistook guns for tools more often after a White prime ($M = .17$, $SD = .14$) than after a Black prime ($M = .12$, $SD = .11$). Simple effects tests showed that the race difference was significant both for tools and for guns, $p\text{s} < .05$. PDP estimates showed that the automatic bias to respond "gun" was higher after Black primes ($M = .62$, $SD = .25$) than after White primes ($M = .54$, $SD = .23$), $F(1, 68) = 5.0$, $p < .05$. In contrast, there was no difference in the cognitive control estimate between Black prime ($M = .68$, $SD = .24$) and White prime conditions ($M = .68$, $SD = .23$), $F(1, 68) < 1$. A parallel pattern of results was expected in the word evaluation task.

As predicted, in the word evaluation task, participants mistakenly responded "bad" more often in the presence of Black faces ($M = .29$. $SD = .12$) than in the presence of White faces, ($M = .24$. $SD = .09$) $F(1, 76) = 26.29$, $p < .001$. They mistakenly responded "good" more often in the presence of White faces ($M = .35$. $SD = .12$) than in the presence of Black faces, ($M = .29$. $SD = .11$), $F(1, 76) = 26.73$, $p < .001$. The Race $\times$ Target word interaction was significant, $F(1, 76) = 42.96$, $p < .001$. PDP estimates showed a highly reliable race difference for automatic bias, $F(1, 76) = 44.88$, $p < .001$. Automatic estimates (scaled such that higher values were more negative) were higher in the presence of Black faces ($M = .50$, $SD = .13$) than in the presence of White faces ($M = .40$, $SD = .11$). In contrast, there were no differences in the cognitive control estimate for Black faces ($M = .42$, $SD = .17$) versus White faces ($M = .41$, $SD = .16$), $F(1, 76) < 1$. These results converge nicely with the results on the weapon identification task. Finally, replicating the typical IAT effect, participants made a higher proportion of errors on the incompatible blocks ($M = .47$, $SD = .08$) than on the compatible blocks ($M = .39$, $SD = .09$), $F(76) = 4.96$, $p < .001$.

Individual Differences

To examine the correlations between individual differences on each measure, I converted the pattern of errors on each task into a single *performance bias* score. The performance bias score was computed as the proportion of stereotype-consistent errors (false "gun" or "bad" responses on Black prime trials + false "tool" or "good" responses on White prime trials) minus the proportion of stereotype-inconsistent errors (false "tool" or "good" responses on Black prime trials,+ false "gun" or "bad" responses on White prime trials). This scoring method is consistent with the method frequently used to score implicit measures (e.g., Wittenbrink et al., 1997). The score represents the effect size of the Prime × Target interaction, reflecting the net stereotypicality of errors; higher values indicate greater stereotyping.

PDP estimates of cognitive control and automatic bias were also converted into a single score for each participant. Because there were no race differences in cognitive control and the cognitive control estimates were highly correlated across Black and White prime conditions ($r = .89$), a single cognitive control value was computed by averaging across both prime conditions. The stereotyping effect in the automatic bias estimate is captured by the difference between Black and White prime conditions. To compute a single automatic estimate for each participant, I partialed out automatic bias scores in the White prime condition from the automatic bias scores in the Black prime condition. The residual scores represented the degree to which participants were more biased to respond in negative ways ("gun" or "bad") in the Black prime condition than in the White prime condition. Parallel analyses using difference scores revealed similar patterns of results. However, the residual method was preferred to the difference score method because difference scores can decrease reliability (Cohen & Cohen, 1983).

Process estimates. Table 1 displays the correlations between the MCPR, antisaccade task, estimates of cognitive control, and estimates of automatic bias. It is worth noting first that for both tasks, automatic and controlled estimates from the same task were uncorrelated (rs $= -.02$ and $.03$). This means that participants with higher control did not have lower automatic stereotype activation; the two were simply independent. The main predictions were that measures reflecting cognitive control would cohere together and those measures reflecting automatic racial bias would cohere together. The theoretically predicted relationships are marked in bold type in the table. Results using the PDP estimates are readily interpretable. Antisaccade performance correlated positively with estimates of cognitive control for the weapon task ($r = .35$, $p < .05$) and the evaluative priming task, although this latter relationship did not reach conventional levels of significance ($r = .19$, $p = .09$). However, because this relationship was theoretically predicted, it was appropriate to use a one-tailed significance test, by which standard the correlation was significant. Estimates of cognitive control from those two tasks correlated significantly with each other ($r = .33$, $p < .05$). This set of relationships is consistent with the idea that the antisaccade task and the cognitive control estimates from the two race bias tasks all draw on similar resources.

To examine the role of motivation to control prejudice, I scored the MCPR results as two subscores: concern with acting prejudiced and restraint to avoid dispute. The Restraint factor did not relate significantly to any of the variables of interest. However, the Concern factor was significantly related to the cognitive control estimates from both the weapon identification task ($r = .25$, $p < .05$) and the evaluative priming task ($r = .33$, $p < .05$). However, the MCPR Concern factor was unrelated to the antisaccade task. This selective pattern of correlations supports the prediction that motivated individuals may exert greater effort at race-related tasks, although they do not show better cognitive control on tasks unrelated to prejudice.

Finally, measures of automatically activated stereotypes were expected to relate to each other. Supporting this prediction, estimates of automatic bias from the weapon and evaluative priming tasks correlated significantly with each other ($r = .29$, $p < .05$). Estimates of automatic bias were also related to the IAT ($r = .17$ for the weapon task, and .19 for the word evaluation task). These latter two correlations were significant by a one-tailed test. Although the correlations are small, this selective pattern of correlations is consistent with the idea that the three prejudice-related tasks all reflect an implicit race bias factor. It is interesting that the IAT was also significantly related to the antisaccade task, $r = -.25$, $p < .05$. Participants with greater attentional control showed less bias on the IAT. This relationship may suggest that the IAT also involves cognitive control processes, consistent with other recent research (e.g., McFarland & Crouch, 2002; Mierke & Klauer, 2001).

The correlational results suggest two distinct sets of relationships. To more formally test whether these measures reflect two independent factors, I conducted a factor analysis on these measures. The factor analysis used oblique rotation so that the correlation between factors could be examined. This analysis showed that two separate factors emerged with eigenvalues greater than 1,

Table 1
Correlations Among Measures of Automatic and Controlled Processes in Experiment 1

	MCPR (concern)	Antisaccade	Control (weapon)	Control (evaluative)	Automatic (weapon)	Automatic (evaluative)
MCPR (concern)	—					
Antisaccade	.04	—				
Control (weapon)	**.25***	**.35***	—			
Control (evaluative)	**.33***	**.19**	**.33***	—		
Automatic (weapon)	.20	−.13	−.02	−.09	—	
Automatic (evaluative)	−.22	−.03	.20	.03	**.29***	—
IAT	.16	−.25*	.11	.06	.17	**.19**

Note. Correlations including the weapon identification task are based on $N = 68$; all other correlations are based on $N = 76$. Theoretically predicted relationships are in bold type. MCPR = Motivation to Control Prejudiced Responses Scale, concern (with acting prejudiced) subscale; Control = cognitive control estimate; Weapon = weapon identification task; Automatic = automatic bias estimate; Evaluative = priming task; IAT = Implicit association task.
* $p < .05$, two-tailed test.

accounting for 52% of the total variance. Table 2 displays the factor loadings. On the first factor loaded the cognitive control estimates from both tasks and antisaccade performance. On the second factor loaded the automatic bias estimates from both tasks and the IAT. Finally, the two factors were not correlated, $r = -.04$. These results provide additional evidence that the estimates of cognitive control and automatic bias reflect distinct factors, each of which was related specifically to other measures of the same construct.

Performance bias. Counter to initial expectations, the correlation between performance bias scores on the weapon task and the evaluative priming task was not significant, $r = -.15$. In addition, performance on the antisaccade task was not significantly related to the performance bias on the weapon task ($r = -.16$) or the evaluative priming task ($r = -.03$). On the face of it, this finding seems to undermine the idea of process generality across tasks. However, it is important to note that the performance bias examined here is the joint product of automatic bias and failures of cognitive control (see Equation 2). The implications of this relationship will be clearer after examining the interactions between automatic and controlled components.

Interactions Between Automatic and Controlled Components

The PDP model specifies a theoretical relationship between automatic bias, cognitive control, and stereotypical performance. Automatic bias is expected to drive stereotype-consistent errors when control fails. Therefore, this model predicts an interaction between automatic bias and cognitive control for stereotype-consistent performance. These predictions were tested using multiple regression analyses, with the performance bias score as a dependent variable. Performance bias was regressed on automatic bias, cognitive control, and the Automatic Bias × Cognitive Control interaction. However, it would be tautological to predict performance on a task from estimates on that same task. Instead, the control estimates from the other tasks were interchanged, yielding four different analyses. The logic behind these analyses is that if estimates of cognitive control go beyond describing an aspect of task performance on a particular task to reveal a more general capacity, then control estimates from one task should moderate the expression of automatic biases on other tasks.

Specifically, in the first analysis (see Figure 1A), stereotypically biased errors on the weapon task were predicted using the automatic estimate from the weapon task and the control estimate from the word evaluation task. In the second analysis (Figure 1B), stereotypically biased errors on the weapon task were predicted using the automatic estimate from the weapon task and the control estimate from the antisaccade task. In the third analysis (Figure 1C), stereotypically biased errors on the word evaluation task were predicted using the automatic estimate from the word evaluation task and the control estimate from the weapon task. In the fourth analysis (Figure 1D), stereotypically biased errors on the word evaluation task were predicted using the automatic estimate from the word evaluation task and the control estimate from the antisaccade task. Thus, each estimate of cognitive control was inserted into the analysis of each other task. The PDP framework predicts a significant interaction between automatic and controlled estimates, such that automatically activated stereotypes are more strongly related to stereotypically biased behavior when cognitive control is low. Figure 1 displays the relationship between automatic estimates and performance bias (the behavioral error index) for participants high (1 *SD* above the mean) and low (1 *SD* below the mean) in cognitive control. All variables were standardized before entering them in regression analyses. In all analyses, automatic estimates were strongly related to performance bias scores (all βs $> .60$). This is to be expected, given that both performance bias scores and PDP automatic estimates depend on stereotypical error rates. The question of interest was how this relationship changed as a function of cognitive control. In each analysis, performance bias scores are the dependent variable.

In the first analysis (Figure 1A) using the control estimate from the word evaluation task, the Automatic × Controlled estimate interaction was significant, $B = -.25$, $t(68) = -2.09$, $p < .05$. In the second analysis (Figure 1B) using the controlled estimate from the antisaccade task, the interaction was not significant, $B = -.03$, $t(68) = -.23$, $p = .82$. In the third analysis (Figure 1C) using the controlled estimate from the weapon task, the interaction was significant, $B = -.17$, $t(68) = -3.23$, $p < .01$. Finally, in the fourth analysis (Figure 1D) using the controlled estimate from the antisaccade task, the interaction was significant, $B = -.12$, $t(76) = -2.35$, $p < .05$. In sum, in three of the four possible combinations of estimates, executive control estimates from other tasks moderated the relationship between automatic bias and behavioral responses. This pattern is important because it suggests a functional equivalence across different measures of executive control beyond the statistical associations revealed by the correlations reported earlier.[2]

Table 2
Factor Analysis of Cognitive Control and Automatic Bias Measures in Experiment 1

Variable	Factor 1 loadings	Factor 2 loadings
Antisaccade	**.72**	−.24
Cognitive control (weapon)	**.82**	.16
Cognitive control (evaluative)	**.62**	−.01
Automatic bias (weapon)	−.18	**.67**
Automatic bias (evaluative)	.13	**.75**
Implicit association task	−.02	**.65**

Note. Weapon = Weapon identification task, evaluative = priming task.

Discussion

Estimates of automatic and controlled processes formed two distinct clusters, supporting the prediction of separate correlates.

[2] A parallel set of analyses was conducted by substituting the automatic bias estimates from each task into the other task. Although the automatic estimates showed the expected pattern of simple correlations (as reported in the main results), none of these analyses showed a significant Automatic × Controlled estimate interaction. One likely reason, as elaborated in the discussion, is that the automatic estimates showed lower reliability than the control estimates.

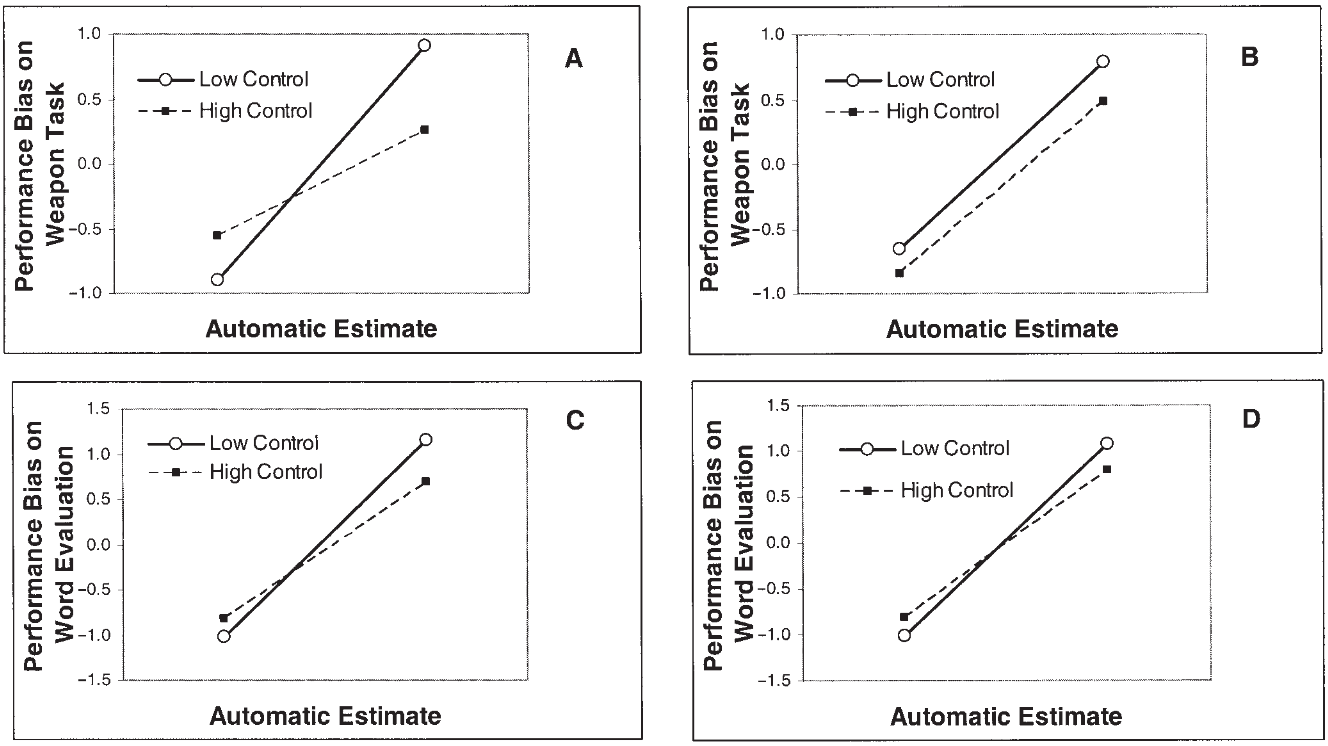

Figure 1. Regression lines showing the relationship between observable errors (performance bias) and automatic stereotyping estimates for participants high and low in cognitive control estimated from various tasks. Panel A shows control estimates from the word evaluation task inserted into analysis of the weapons task. Panel B shows control estimates from the antisaccade task inserted into the analysis of the weapon task. Panel C shows control estimates from the weapon task inserted into analysis of the word evaluation task. Panel D shows control estimates from the antisaccade task inserted into analyses of the word evaluation task.

Measures of cognitive control formed one coherent factor, whereas measures of automatic racial bias formed a second factor. This pattern of results supports the hypotheses that PDP estimates in the weapon task reflect automatically activated racial associations, on one hand, and executive control on the other. The degree of executive control exerted on the weapon identification and word evaluation tasks appeared to have been related to both general executive ability and to race-specific motivations. Supporting the hypothesized relationship in three of four tests, cognitive control estimates from each task moderated the expression of automatic bias in overt performance.

An implication of these findings is that the relationship between overt stereotypical errors and automatic bias depends on how much control a person exerts. Consulting Equation 2, it can be seen that at one extreme, for a person with a score of 0 for control, stereotypical errors are a one-to-one reflection of automatic bias. At the other extreme, for a person with perfect control, stereotypical errors are completely uninformative about automatic bias (in fact, stereotypical errors/incongruent must equal 0). Along the middle range of the continuum of control, the more control a person has, the more his or her automatic bias is underestimated with the performance bias score. This relationship has important implications for explaining why the PDP estimates correlated systematically, even though the performance scores did not.

Two individuals with identical automatic biases can have different performance bias scores because of differences in control. In fact, if two individuals differ in automatic bias (e.g., one has an automatic bias of .55 and the other of .60), their rank orders can actually *reverse* on the performance bias score if the second person has much better control than the first. Because correlations depend on rank orders, this potential flipping of rank orders may be an important factor in explaining why performance scores did not correlate as predicted.

A second reason may be reliability of measurement. Supplementary analyses showed that the process estimates were more reliable than performance bias scores. To calculate reliability, I calculated scores for each of the four blocks of the weapon identification and word evaluation tasks. The scores from each block were used to calculate Cronbach's alpha as if a 4-item scale were being used. Consistent with other findings with priming measures, performance scores showed very low reliability (weapon identification, α = .12; word evaluation, α = .28). Automatic estimates showed somewhat better reliability (weapon identification, α = .33; word evaluation, α = .55) although these values were still quite low. Finally, control estimates showed much better reliability (weapon identification, α = .89; word evaluation, α = .84). Because low reliability attenuates correlations, this probably plays an important role in explaining why the PDP estimates showed more systematic relations than performance bias scores. Leaving automatic and controlled components unmodeled can both alter rank orders and reduce reliability. Beyond the present results, this explanation may shed light on why other studies have found that many implicit measures are high in measurement error and tend to be poorly correlated with each other (Cunningham et al., 2001).

Low reliability may also explain why the relationships observed here, although reliable and systematic, were small in magnitude. A latent variable approach that controls for measurement error may reveal that the relationships between latent constructs are in fact more substantial than the relationships between observed variables (Cunningham et al., 2001). Unfortunately, latent variable analyses generally require larger sample sizes than are available in the present study. A second implication of the small correlations may be that even though different conceptually related process estimates draw on overlapping cognitive processes, they are likely to have substantial unique, or unshared, components as well. This possibility is consistent with theorizing about the nature of executive control as a suite of interrelated processes rather than a monolithic concept and is also consistent with empirical studies that have repeatedly found that different measures of executive functioning have both shared and unique components (e.g., Friedman & Miyake, 2004).

A potential weakness of the present results is the fact that all of these measures were somewhat similar cognitive tasks. Each task required responding to a stimulus on the computer monitor. The three race-related tasks, but not the antisaccade task, required quick responding. One might argue that the relationships observed were, in part or in whole, driven by individual differences in skill at computer-based tasks. It would be difficult to use this argument to explain the selective relations among automatic bias estimates and the relationship between motivations to control prejudice and cognitive control estimates, but it could at least account for the relations between control estimates. In any case, it would be useful to gather evidence that the processes estimated by the PDP estimates related to other social judgments beyond detection tasks. This goal was taken up in Experiment 2, which examined whether automatic bias and cognitive control estimates could predict race-biased judgments of other people.

Experiment 2: Impression Formation

Central executive functions are considered "central" because they coordinate and regulate virtually all kinds of information processing. Such a fundamental set of processes should be expected to have pervasive effects (e.g., Feldman Barrett, Tugade, & Engle, 2004). Impression formation was chosen to extend the generality of the present analysis. This topic has traditionally been studied using sequential correction or deliberative reasoning models of control. If executive control plays an important role in this domain, it would suggest that social judgments far removed from cognitive interference tasks are influenced by the interplay of automatic activation and executive control in the manner described in the first experiment.

We used a behavioral description of a target person named "John." This description included individuating behavioral information, but the information was ambiguous on the dimensions of intelligence and aggressiveness (Higgins, Rholes, & Jones, 1977; Srull & Wyer, 1980). Pilot testing showed that, holding all else equal, people tend to rate John as moderate on both traits. Previous research has also shown that when John is described as Black, people with more negative attitudes toward Blacks in general (measured either explicitly or implicitly) tend to rate John more negatively (Lambert, Cronen, Chasteen, & Lickel, 1996; Lambert, Payne, Ramsey, & Shaffer, 2004). Our previous research used two different reaction time–based implicit measures, which did not permit examination of process estimates. The hypothesis of Experiment 2 was that executive control (as measured by PDP estimates) would moderate that relationship.

Individuals exerting greater executive control are better able to resist interference from irrelevant information, constraining their processing to the most relevant inputs. They are also better able to make strategic responses in novel, uncertain, or ambiguous circumstances (Norman & Shallice, 1986). In forming impressions of a new (Black) person, automatically activated attitudes and stereotypes may be a source of interfering information that biases judgment (von Hippel et al., 2000). When people are not exerting much executive control, their judgments are likely to be affected by these automatically activated sources of information. However, when people are engaging more extensive control processes, they may be less likely to incorporate these attitudes and stereotypes into their judgments, instead forming more strategic judgments under uncertainty. Therefore, it was predicted that the most negative impressions of John would be formed by those with a strong automatic racial bias *and* poor executive control.

Method

Participants

Fifty-six non-Black participants (37 women, 19 men) took part in the experiment for course credit. One participant was excluded for not following directions and, hence, producing a below-chance performance.

Procedure

After providing informed consent, participants completed the weapon identification task. This task was identical to the one used in Experiment 1, with the exceptions that the response deadline was 550 ms rather than 500 ms to reduce the number of missed responses and the number of trials was increased to 320 to increase reliability. After completing the weapon identification task, participants were asked to take part in a second, ostensibly unrelated study: the impression formation task. The experimenter told participants that we were interested in how people form first impressions of other people. Participants were presented with a biographical sketch based on one used in the study by Lambert et al. (1996) in which the target person had supposedly filled out his name, gender, address, place of birth, educational status, academic major, expected graduation date, citizenship status, and an identification number, as well as his racial/ethnic background (which was checked "African American"). The target's race was thus only one of a dozen or so pieces of background information presented about the target person. In earlier studies, we had manipulated the race of the target character. Both explicitly and implicitly measured attitudes correlated with judgments when the target was described as Black but not when the target was identified as White (Lambert et al., 2004). That is, there was no overall "harshness" effect associated with racial attitudes. Therefore, the present study designated the character as Black in all cases in order to maximize statistical power. The text of the passage was as follows:

> After working for a while, John looked up from his books to have another look at the letter that had been sitting on his desk. He had gotten a 2.7 again and hadn't gotten any As for the third semester in a row. His parents were just a bit worried about his grades. If things went the way they had been going, it looked as if he was going to get mostly Bs this semester, with maybe a few Cs. His performance so far made him a bit down, but he resolved to do better in the future. It was the only way that he'd be able to get into that graduate school on the East Coast that he had heard so much about.
>
> After his morning classes, John grabbed some lunch at the cafeteria. The place was a little crowded, but John found a table in the back and sat down. He thought about how much he was looking forward to going home. John thought how nice it would be to eat some real food instead of the tasteless stuff they served at school.
>
> Later on that day, John needed to do a couple of important errands in the city, but unfortunately his car started making some noises. John thought it might be something pretty serious, so he looked for a shop that could fix it. John found a shop and when he talked to the mechanic, John told him that he would have to go somewhere else if he couldn't fix his car that same day. While he was waiting for the car to be fixed, John went to a store and demanded his money back from the sales clerk for something he had bought earlier in the week.

Participants reported their overall evaluation of the target on a scale ranging from −5 (*very unfavorable*) to +5 (*very favorable*) and indicated how much they would want to meet this person along a scale ranging from −5 (*wouldn't want to meet him*) to +5 (*would want to meet him*). After this, they estimated the degree to which the target possessed a number of specific traits (*likable*, *successful*, *unfriendly*, *intelligent*, *competent*, *unmotivated*, *patient*, *self-assured*, *incompetent*, *polite*, *lazy*, *bright*, *argumentative*, *aggressive*, *hardworking*, *athletic*, *easy to get along with*, *cooperative*, *hostile*, *shy*, *responsible*, and *ambitious*). These judgments were made on a scale ranging from 0 (*not at all*) to 10 (*extremely*).

Following the impression formation task, participants completed some ancillary measures, including the Modern Racism Scale (McConahay, 1986) and the MCPR. Participants received an educational debriefing in which the purpose and hypotheses of the experiment were explained.

Results

Mean Results for Weapon Identification Task in Experiment 2

A 2 (Prime Race) × 2 (Target Object) ANOVA was used to analyze the results of the weapon identification task, with errors as the dependent variable. The predicted Prime Race × Target Object interaction was significant, $F(1, 54) = 10.15$, $p < .01$. Participants mistook tools for guns more often after a Black prime than after a White prime (*M*s = .12 vs. .10) and mistook guns for tools more often after a White prime than after a Black prime (*M*s = .10 vs. .09). Simple effects tests showed that the race difference was significant for false "gun" responses, $F(1, 54) = 5.85$, $p < .05$, and marginally significant for false "tool" responses, $F(1, 54) = 3.52$, $p = .07$. PDP estimates showed that the automatic bias to respond "gun" was higher after Black primes ($M = .42$) than after White primes ($M = .35$), $F(1, 54) = 8.47$, $p < .01$. In contrast, no difference was seen in the cognitive control estimate between Black prime ($M = .73$) and White prime conditions ($M = .74$), $F(1, 54) < 1$. These findings replicate the typical pattern observed in mean performance. However, the main focus of this study was on individual differences, which are discussed next.

Correlations With Impression Ratings in Experiment 2

To reduce the set of trait ratings to a simpler set of meaningful dimensions, I conducted a factor analysis. This analysis used oblique rotation, which allowed related components to correlate with each other. The analysis produced a total of six factors that accounted for 74% of the variance. However, the first component explained 45% of the total variance and was clearly the dominant and most easily interpretable factor. This component appeared to represent a general evaluative dimension and contained high loadings on many of the individual items.

The items loading above .50 for this factor were as follows: *trustworthy* (.83), *intelligent* (.83), *cooperative* (.83), *overall evaluation* (.82), *hardworking* (.81), *competent* (.81), *ambitious* (.80), *polite* (.79), *likable* (.79), *responsible* (.77), *would like to meet* (.76), *successful* (.76), *bright* (.73), *easy to get along with* (.71), *patient* (.66), *energetic* (.66), *incompetent* (–.66), *impolite* (–.65), *unmotivated* (–.64), *hostile* (–.60), *self-assured* (–.57), *impatient* (–.57), *unfriendly* (–.53), and *argumentative* (–.53). A weighted index capturing the pattern of loadings for this component was used in subsequent analyses. To be consistent with the scaling of other measures, I scaled the index so that higher values represented a more *negative* impression of the target character.[3]

First, the performance bias score from the weapon identification task was computed. As in Experiment 1, higher scores represent a stronger Prime Race × Target Object interaction or a more stereotypical error pattern. The performance bias score was significantly correlated with impressions of the target character, $r = .40$, $p < .01$. Participants with a more stereotypical pattern of weapon misidentifications formed a more negative impression of the Black target character.

The most obvious interpretation of this finding is that the performance bias score represents implicit negative attitudes toward Blacks. When confronted with ambiguous behaviors by a Black individual, participants with more implicit negativity disliked him more. However, as argued in the introduction and demonstrated in both present studies, the performance score represents the joint influence of automatic racial bias and executive control. To examine the separate contributions of each factor, I used a regression analysis to analyze PDP estimates. Figure 2 shows perceptions of the target character plotted as a function of automatic bias, at 1 *SD* above and 1 *SD* below the mean of cognitive control. The regression analysis confirmed that automatic bias, cognitive control, and their interaction all contributed

[3] To examine impressions on specifically stereotype-relevant dimensions, I computed two other indexes based on the stereotypical dimensions of intelligence and aggressiveness. These indexes showed the same general pattern of relationships as the global evaluation reported, although they were statistically weaker individually. This is likely because these indexes included fewer items, rendering less reliable measures. In any case, stereotypic ratings did not reveal any qualitative differences from the evaluative ratings.

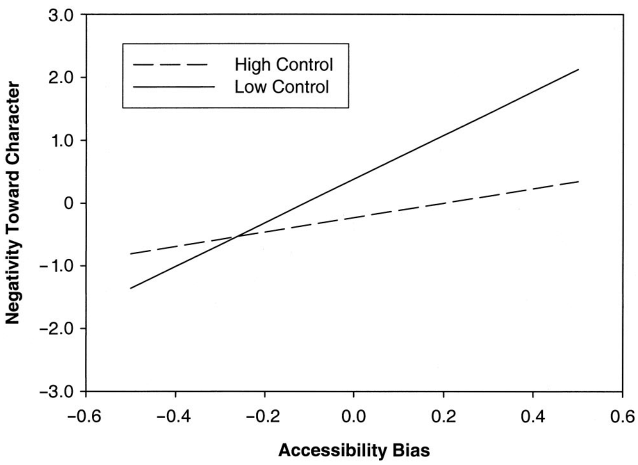

Figure 2. Regression lines predicting impression ratings from automatic bias for participants high and low in cognitive control in Experiment 2.

to people's judgments. People with stronger automatic bias formed more negative impressions of the target character, $\beta = .34$, $t(54) = 2.72$, $p < .01$. In addition, people with poorer cognitive control formed more negative impressions of the target character, $\beta = -.33$, $t(54) = 2.66$, $p < .01$. Finally, the interaction between automatic bias and cognitive control was significant, $\beta = -2.17$, $t(54) = 2.03$, $p < .05$ ($F_{\text{change}} = 3.99$, $p < .05$). The relationship between automatic bias and target judgments depended on the degree of cognitive control. Supporting the main prediction, the relationship between automatic bias and target judgments was stronger among participants with poor cognitive control.[4]

As a final analysis, the reliability of process estimates was examined. The performance bias score in this experiment showed better reliability than in the first experiment, most likely because of the larger number of trials ($\alpha = .60$). The automatic estimate showed much better reliability ($\alpha = .92$), and the control estimate also showed high reliability ($\alpha = .91$). Averaging across studies (weighted by sample size), the reliability was $\alpha = .32$ for performance bias, $\alpha = .59$ for the automatic estimate, and $\alpha = .88$ for the control estimate. In both studies, PDP estimates were more reliable than performance bias scores alone.

Discussion

Participants with a strong automatic bias formed a more negative impression of the target character, but this was especially so if they were also low in cognitive control. Uncovering the joint role of automatic bias and control processes in such judgments offers several novel theoretical contributions. First, this relationship provides additional evidence that the process estimates have utility beyond concretely measuring visual discriminations in the weapon identification task. The estimates appear to measure consequential psychological properties that are predictive of other fairly removed social judgments.

A second insight is that when task performance was distilled into separate process estimates, the estimates provided more information than traditional performance bias scores. Researchers in the relatively young but fast-growing field of implicit attitude measurement have recently been intensely interested in the predictive validity of implicit measures. The present results suggest that in some cases, the predictive validity of a task may be underestimated by performance scores alone, especially among participants who are high in cognitive control.

A third contribution of this framework is that it provides is a straightforward measurement model for quantifying estimates of the underlying processes. This measurement model can be readily implemented in many paradigms currently being used, provided that there are enough errors for meaningful analysis. Consequently, the present approach may facilitate theory testing by helping to directly compare different process-oriented theoretical explanations for interesting effects.

The demonstrated interaction between automatic and controlled processes in predicting overt behaviors and judgments might suggest a trivializing alternative explanation. Because the process estimates are derived from error performance using equations that imply such an interaction, is it mathematically necessary that any variable related to the degree of bias in actual error rates also be related to the Automatic × Controlled interaction? If so, then demonstrating the Automatic × Controlled interaction might be tautological and provide no new insights.

However, several of the dissociations that have been found using the weapon identification paradigm argue against this explanation. These dissociations showed that some variables, such as speeded responding, increased stereotypical errors by reducing control without changing automatic bias (Payne, 2001). Other variables, such as warning about the effects of stereotypes, increased stereotypical errors by increasing automatic bias without changing control (Payne et al., 2002). In both of these studies, the manipulation increased the net stereotypicality of errors, which is identical to the performance bias scores I used. This net bias can be changed in three different ways:

1. a main effect change in automatic bias,

2. a main effect change in cognitive control, or

[4] The Modern Racism Scale (MRS; McConahay, 1986) and the Motivation to Control Prejudiced Responses Scale (MCPR) were included in this study after all other measures had been collected. Because the results are of secondary interest for the hypotheses tested, these results are only briefly reported here. The MRS showed a small relationship to impression judgments, $r = .24$, $p < .10$, but not to any other measures of interest. Replicating the general pattern shown in Experiment 1, the Concern subscale of the MCPR correlated positively with the process dissociation procedure control estimate, although this relationship did not reach significance in this study, $r = .17$, $p = .23$. This is likely due, in part, to the smaller sample size in the second study. A second reason may be that participants in this experiment were only exposed to one task that was obviously race-related. In Experiment 1, participants engaged in three obvious race-relevant tasks. This protocol may have more strongly prodded participants who were motivated to avoid prejudice to exert greater effort at the tasks in Experiment 1. Also consistent with Experiment 1, the restraint subscale did not correlate with the control estimate, $r = .01$. Finally, a regression analysis controlling for both MRS and MCPR scores showed no significant changes in the relationships between automatic and control estimates and impression judgments, suggesting that these relationships emerged independent of explicit attitudes and motivations.

3. a change in both automatic and control components above and beyond these main effects, represented by the Automatic × Control interaction.

To empirically confirm that the performance bias index is not necessarily identical to the Automatic × Control interaction, I reanalyzed the results from these two experiments. An analysis of covariance was used to test whether the experimental manipulation affected the Automatic × Control interaction after controlling for the main effects of automatic and control alone. To do this, I calculated an interaction term by multiplying automatic and control estimates, after centering them. The rationale is the same as that used when testing interaction effects in regression models. In both of these studies, the manipulation significantly affected the performance bias score. Speeded responding selectively reduced the control estimate, whereas warning about racial stereotypes significantly increased the automatic estimate. However, when the main effects of automatic and control were controlled, neither manipulation significantly affected the Automatic × Control interaction term. This pair of analyses shows that variables can be related to the performance bias score without necessarily being related to the Automatic × Control interaction. Stereotypical errors may be related to changes in automatic bias alone, or to changes in cognitive control alone, or to the interactive relationship of both, as demonstrated by the present studies. Only by modeling the underlying processes can these different possibilities be empirically disentangled.

General Discussion

The present studies highlight the role that individual differences in executive control play in moderating automatic social biases. Individuals with good executive control showed the same level of automatic stereotype activation as those with poor control. However, automatically activated stereotypes were less likely to be expressed in behavioral errors or social judgments among those with good executive control. These studies demonstrate that executive control can exert a broad influence on social judgment, from performance on an implicit stereotyping task to impression formation.

Although executive control has become a critical component of research and theory in many areas of cognitive science, it has been less fully integrated in social psychological theories (even in dual-process theories whose central topic is the interaction of automatic and controlled processes). However, this may be changing. For instance, in a theoretical review applying executive control to dual-process theories, Feldman Barrett et al. (2004) proposed that individuals high in executive control "should be better able to use symbolically represented rules to monitor decisions and behaviors" (p. 561). They reasoned that this might lead individuals with high control to better regulate the influence of activated stereotypes on judgments. The present results are entirely consistent with this prediction. The behavior of high-control individuals was in fact less influenced by their own activated stereotypes. It should be noted that this is different than saying that high-control participants experienced less automatic stereotype activation; they did not. It is also different from personal struggles to overcome stereotypic thoughts (Devine, 1989). Instead, the ability to focus attention on relevant inputs while avoiding interference served as a sort of gatekeeper between automatic activation and overt discrimination.

Beyond documenting the role of executive control as a gatekeeper, the present studies provide evidence for a measurement model that is useful for quantifying automatic and controlled processes. The main conclusions in the present research depend on application of the PDP. Because this procedure depends on a number of assumptions, some of which have proven controversial, it is important to consider whether the assumptions are justified. The following section considers two important assumptions. The first is that automatic and controlled processes exert independent influences on responses. The second is that automatic and controlled estimates represent the same processes in the congruent and incongruent experimental conditions. For a more extended discussion of the assumptions and boundary conditions used in the PDP of memory research, the reader is referred to Curran and Hintzman (1995); Jacoby and Shrout (1997); and Jacoby, Begg, and Toth (1997).

Assumptions Underlying the PDP

The most controversial assumption underlying the PDP is that automatic and controlled processes exert independent influences on performance. The most commonly used method for validating the independence assumption is to show dissociations, or selective effects, on the two process estimates (Jacoby et al., 1997). The logic behind this approach is that if automatic and controlled processes are independent, then one should be able to influence one process without influencing the other. However, if the independence assumption is badly violated and thus automatic and controlled processes are strongly related, then any manipulation influencing one process is likely to also affect the other.

Several experiments have demonstrated dissociations between automatic and controlled processes in the context of stereotyping tasks. In repeated tests, race primes influenced the automatic estimate but not the controlled estimate (Amodio et al., 2004; Lambert et al., 2003; Payne, 2001; Payne et al., 2002). At the same time, fast responding reduced the controlled estimate but had no effect on the automatic estimate (Payne, 2001; Payne et al., 2002). Lambert et al. (2003) found that anxiety over a public discussion affected the controlled estimate but had no effect on the automatic estimate. Amodio et al. (2004) documented selective correlations between controlled estimates and ERP signals, reflecting conflict monitoring. Govorun and Payne (2005) showed that ego depletion reduced the controlled estimate but not the automatic estimate. Finally, the present studies demonstrate distinct correlates of the automatic and controlled estimates. Each of these studies provides evidence that automatic and controlled processes may be influenced independently of each other, supporting the independence assumption as applied to the weapon identification paradigm.

The in-concert (congruent) and in-opposition (incongruent) conditions in the weapon identification task parallel inclusion and exclusion instructions in other PDP studies (e.g., Jacoby, 1991). An important assumption with this method is that automatic and controlled processes play the same roles in inclusion and exclusion conditions. This assumption can be seen in the PDP equations by noting that cognitive control and automatic bias are compared across congruent and incongruent conditions. Therefore, controlled in the congruent condition must mean the same thing as controlled in the incongruent condition; automatic in the congruent condition must mean the same thing as automatic in the incongruent condition. Although the original applications of PDP used inclusion and exclusion instructions, alternative forms of the procedure have been developed that avoid use of the different instructional sets (Hay & Jacoby, 1996). These alternative methods rely on congruence or incongruence between automatic and controlled processes rather than differing instructions. An advantage of manipulating congruence rather than instructions is that participants do not have to shift instructional set, which can potentially change the meaning of processes across inclusion and exclusion conditions (Graf & Komatsu, 1994). The weapon identification task manipulates congruence rather than instructions, minimizing the likelihood that the meaning of processes varies across congruent and incongruent trials because of differences in interpretation. Moreover, if the role of automatic and controlled parameters differed substantially between congruent and incongruent conditions, it would be difficult to obtain the clean dissociations reviewed above.

Comparison of Dual-Process and Single-Process Models

The preceding section laid out some of the assumptions that must be met to properly apply the PDP. An explicit statement of those assumptions may lead some readers to prefer other models in an effort to avoid such assumptions, but that view overlooks the fact that all formal models, like all statistical tests, rely on mathematical assumptions. This section compares the dual-process model of PDP to the single-process model represented by SDT Tanner & Swets, 1954). SDT is the most closely related alternative model and has been invoked to explain results such as stereotypical weapon misidentification. I argue that the assumptions of SDT are not fewer than, but simply different from, those of PDP. Further, both SDT and PDP analyses require theoretical assumptions about the nature of underlying information processing. These theoretical assumptions, even if not explicitly acknowledged, constrain the plausible interpretations of the results of each model.

SDT has been used to analyze results in stereotyping tasks similar to the weapon identification task (e.g., Correll, Park, Judd, & Wittenbrink, 2002; Greenwald, Oakes, & Hoffman, 2003). Analyses using SDT and PDP often yield similar results, with SDT estimates of sensitivity related to PDP estimates of cognitive control and SDT estimates of criterion related to PDP estimates of automatic bias (Greenwald et al., 2003). However, the estimates produced by these models afford very different interpretations. SDT is a single-process model in the sense that a single dimension of information is processed. According to SDT, discriminations are made on the basis of a single continuum of evidence strength. For simple psychophysics tasks, the evidence may be the strength of a perceptual signal. For memory tasks, evidence strength represents familiarity or strength of a "memory trace." Participants must distinguish between a distribution of evidence strength generated by nontargets (which represents noise) and a distribution of evidence strength generated by targets (which includes both signal and noise). Participants accomplish this discrimination by choosing a criterion, or a point on the continuum of evidence strength, that marks how much evidence they require to decide that the target is present. The process is similar to (and the theory was influenced by) the use of statistical decision rules in null-hypothesis testing.

SDT depends on a number of assumptions (see MacMillan & Creelman, 1991; Snodgrass & Corwin, 1988). The first assumption is that the distributions of noise and signal-plus-noise are normally distributed. The second is that the variances of these two distributions are equal. If either of these assumptions is violated, estimates of sensitivity and criterion may be biased. A third assumption also depends on these first two: that sensitivity and criterion are independent of each other. This is the same as the independence assumption in PDP. If the normality assumption or the equal variances assumption is violated, the estimates may not be independent (MacMillan & Creelman, 1991).

In addition to these mathematical assumptions, both SDT and PDP depend on several (different) theoretical assumptions about the cognitive processes underlying performance. In SDT, it is assumed that both sensitivity and criterion estimates are generated from the processing of a single stream of information (i.e., evidence strength). For the weapon identification task and similar paradigms, it is not clear how to best characterize the evidence continuum. Correll et al. (2002) suggested that the dimension could be called "perceived threat." This view treats the processing of gunlike features (e.g., triggers, barrels) and stereotypes about Black individuals alike. Either a Black individual adds to the perceived threat in the same way that the perception of target features does, or participants need less perceived threat from the target features in the presence of a Black individual. Correll et al. (2002) discussed several possible psychological interpretations that could be consistent with their SDT results, including perception of different features, different interpretations of ambiguous features, and differences in decision criteria. In any case, applying SDT requires the assumption that racial stereotypes and features of the target items influence responses through a single dimension of evidence. SDT does not address automatic versus controlled aspects of processing.

With PDP, the assumptions are quite different. As applied to the weapon identification task, the assumption is that the automatic component represents associations activated automatically by racial cues that make up one basis for responding. A second basis for responding is a controlled decision reached by constraining processing to the relevant input. This interpretation of the automatic estimate as reflecting automatic stereotyping and the controlled estimate as reflecting cognitive control has led to the predictions of selective effects just reviewed. In those studies, factors known to affect controlled processing and those known to affect automatic processing caused systematic dissociations. In the present study, people with strong automatic bias *and* poor cognitive control formed the most negative impression of a Black target character. These findings follow naturally from the dual-process account of PDP. It is less clear how an SDT framework would account for these results. For example, being neutral to issues of automaticity and control, there would seem to be little reason to predict that variables affecting controlled or automatic processing should selectively influence one parameter or the other. It may be that some version of SDT could explain these findings, but further theoretical links would need to be elaborated to do so.

Some versions of SDT are more compatible with PDP than the standard single-process version. Although it is not commonly used, multidimensional SDT has been developed to treat situations in which discriminations are based on two or more dimensions of evidence (MacMillan & Creelman, 1991). A two-dimensional SDT model in which one dimension reflected automatic associations and the second dimension reflected controlled processing of target items would be consistent with many aspects of the PDP analysis. Acknowledging the similarities and differences between models is important because each entails different theoretical commitments of which researchers should be aware. However, the dividends of such analyses are considerable. The present framework, as just one example, includes several novel predictions including the dissociation patterns described. This framework also draws links between social cognition and several other fields and different levels of analysis (Amodio et al., 2004; Payne et al., 2004). Despite intense interest in the distinctions between automatic and controlled processes, social psychological models have until recently given little emphasis to executive control. Research on automatic processing in social psychology was facilitated when researchers recognized that automaticity was not monolithic but was a description for several different qualities of information processing (Bargh, 1989). Much less attention has been devoted to understanding the different forms that controlled processing might take, a task worthy of future research.

References

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

Amodio, D. M., Harmon-Jones, E., Devine, P. G., Curtin, J. J., Hartley, S. L., & Covert, A. E. (2004). Neural signals for the detection of unintentional race bias. *Psychological Science, 15,* 225–232.

Baddeley, A. D. (1986). *Working memory.* London: Oxford University Press.

Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influences in social perception and cognition. In J. S. Uleman & J. A. Bargh (Eds.) *Unintended thought* (pp. 3–51). New York: Guilford.

Bodenhausen, G. V. (1990). Stereotypes as judgmental heuristics: Evidence of circadian variations in discrimination. *Psychological Science, 1,* 319–322.

Bodenhausen, G. V., & Macrae, C. N. (1998). Stereotype activation and inhibition. In R. S. Wyer, Jr. (Ed.), *Advances in social cognition* (Vol. 11, pp. 1–52). Mahwah, NJ: Erlbaum.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review, 108,* 624–652.

Botvinick, M. M., Nystrom, L. E., Fissell, K., Carter, C. S., & Cohen, J. D. (1999). Conflict monitoring versus selection-for-action in anterior cingulated cortex. *Nature, 402,* 179–181.

Chaiken, S., & Trope, Y. (Eds.). (1998). *Dual-process theories in social psychology.* New York: Guilford.

Cohen, J., & Cohen, P. (1983). *Applied multiple regression: Correlation analysis for the behavioral sciences.* Hillsdale, NJ: Erlbaum.

Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2002). The police officer's dilemma: Using race to disambiguate potentially threatening individuals. *Journal of Personality and Social Psychology, 83,* 1314–1329.

Cunningham, W. A., Johnson, M. K., Raye, C. L., Gatenby, J. C., Gore, J. C., & Banaji, M. R. (2004). Separable neural components in the processing of Black and White faces. *Psychological Science, 15,* 806–813.

Cunningham, W. A., Preacher, K. J., & Banaji, M. R. (2001). Implicit attitude measures: Consistency, stability, and convergent validity. *Psychological Science, 12,* 163–170.

Curran, T., & Hintzman, D. L. (1995). Violations of the independence assumption in process dissociation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 531–547.

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

Devine, P. G., & Monteith, M. J. (1999). Automaticity and control in stereotyping. In S. Chaiken & Y. Trope (Eds.) *Dual-process theories in social psychology*. New York: Guilford.

Devine, P. G., Monteith, M. J., Zuwerink, J. R., & Elliot, A. J. (1991). Prejudice with and without compunction. *Journal of Personality and Social Psychology, 60,* 817–830.

Dovidio, J. F., Brigham, J. C., Johnson, B. T., & Gaertner, S. L. (1996). Stereotyping, prejudice, and discrimination: Another look. In C. N. Macrae, C. Stangor, & M. Hewstone (Eds.), *Stereotypes and stereotyping.* New York: Guilford.

Draine, S. C., & Greenwald, A. G. (1998). Replicable unconscious semantic priming. *Journal of Experimental Psychology: General, 127,* 286–303.

Dunton, B. C., & Fazio, R. H., (1997). An individual difference measure of motivation to control prejudiced reactions. *Personality and Social Psychology Bulletin,* 23, 316–326.

Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception and Psychophysics, 16,* 143–149.

Everling, S., & Fischer, B. (1998). The antisaccade: A review of basic research and clinical studies. *Neuropsychologia, 36,* 885–899.

Fazio, R. H. (1990). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 23, pp. 75–109). New York: Academic.

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.

Feldman Barrett, L., Tugade, M. M., & Engle, R. W. (2004). Individual differences in working memory capacity and dual-process theories of the mind. *Psychological Bulletin, 130,* 553–573.

Friedman, N. P., & Miyake, A. (2004). The relations among inhibition and interference control functions: A latent-variable analysis. *Journal of Experimental Psychology: General, 133,* 101–135.

Gilbert, D. T., & Hixon, J. G. (1991). The trouble of thinking: Activation and application of stereotypic beliefs. *Journal of Personality and Social Psychology, 60,* 509–517.

Govorun, O. & Payne, B. K. (2005). *From depletion to discrimination: The role of self-control in overriding automatic stereotyping.* Manuscript submitted for publication.

Graf, P., & Komatsu, S. (1994). Process dissociation procedure: Handle with caution. *European Journal of Cognitive Psychology, 6,* 113–129.

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

Greenwald, A. G., Oakes, M. A., & Hoffman, H. G. (2003). Targets of discrimination: Effects of race on responses to weapon holders. *Journal of Experimental Social Psychology, 39,* 399–405.

Hasher, L., & Zacks, R. T. (1988). Working memory, comprehension, and aging: A review and a new view. In G. H. Bower (Ed.) *The psychology* *of learning and motivation: advances in research and theory* (Vol. 22, pp. 193–225). San Diego, CA: Academic.

Hay, J. F., & Jacoby, L. L. (1996). Separating habit and recollection: Memory slips, process dissociations, and probability matching. *Journal of Experimental Social Psychology: Learning, Memory, and Cognition, 22,* 1323–1335.

Hermans, D., De Houwer, J., & Eelen, P. (2001). A time course analysis of the affective priming effect. *Cognition & Emotion, 15,* 143–165.

Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology, 13,* 141–154.

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30,* 513–541.

Jacoby, L. L., Begg, I. M., & Toth, J. P. (1997). In defense of functional independence: Violations of assumptions underlying the process–dissociation procedure? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 484–495.

Jacoby, L. L., & Shrout, P. E. (1997). Toward a psychometric analysis of violations of the independence assumption in process dissociation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 505–510.

Judd, C. M., Blair, I. V., & Chapleau, K. M. (2004). Automatic stereotypes versus automatic prejudice: Sorting out the possibilities in the Payne (2002) weapon paradigm. *Journal of Experimental Social Psychology, 40,* 75–81.

Kane, M. J., Bleckley, M. K., Conway, A. R. A., & Engle, R. W. (2001). A controlled-attention view of working memory capacity. *Journal of Experimental Psychology: General, 130,* 169–183.

Kahneman, D. (1973). *Attention and effort*. Englewood Cliffs, NJ: Prentice-Hall.

Lambert, A. J., Cronen, S., Chasteen, A. L., & Lickel, B. (1996). Private vs. public expressions of racial prejudice. *Journal of Experimental Social Psychology, 32,* 437–459.

Lambert, A. J., Payne, B. K., Jacoby, L. L., Shaffer, L. M., Chasteen, A. L., & Khan, S. K. (2003). Stereotypes as dominant responses: On the "social facilitation" of prejudice in anticipated public contexts. *Journal of Personality and Social Psychology, 84,* 277–295.

Lambert, A. J., Payne, B. K., Ramsey, S., & Shaffer, L. M. (2004). On the predictive validity of implicit attitude measures: The moderating effect of perceived group variability. *Journal of Experimental Social Psychology, 40,* 708–722.

Logan, G. D. (1997). The automaticity of academic life: Unconscious applications of an implicit theory. In R. S. Wyer (Ed.), *Advances in social cognition* (Vol. 10, pp. 157–179). Mahwah, NJ: Erlbaum.

MacMillan, N. A., & Creelman, C. (1991). *Detection theory: A user's guide.* New York: Cambridge University Press.

Macrae, C. N., Bodenhausen, G. V., Schloerscheidt, A. M., & Milne, A. B. (1999). Tales of the unexpected: Executive function and person perception. *Journal of Personality and Social Psychology, 76,* 200–213.

Macrae, C. N., Milne, A. B., & Bodenhausen, G. V. (1994). Stereotypes as energy-saving devices: A peek inside the cognitive toolbox. *Journal of Personality and Social Psychology, 66,* 37–47.

McFarland, S. G., & Crouch, Z. (2002). A cognitive skill confound on the implicit association test. *Social Cognition, 20,* 483–510.

McConahay, J. B. (1986). Modern racism, ambivalence, and the modern racism scale. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 91–126). New York: Academic.

Mierke, J., & Klauer, K. C. (2001). Implicit association measurement with the IAT: Evidence for the effects of executive control processes. *Zeitschrift fur Experimentelle Psychologie, 48,* 107–122.

Musch, J., & Klauer, K. C. (2001). Locational uncertainty moderates affective congruency effects in the evaluative decision task. *Cognition & Emotion, 15,* 167–188.

Myrdal, G. (1944). *An American dilemma: The Negro problem and modern democracy.* New York: Harper & Bros.

Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation: Advances in research and theory (*Vol. 4, pp. 1–18). New York: Plenum.

Otten, S., & Wentura, D. (1999). About the impact of automaticity in the Minimal Group Paradigm: Evidence from affective priming tasks. *European Journal of Social Psychology, 29,* 1049–1071.

Payne, B. K. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. *Journal of Personality and Social Psychology, 81,* 181–192.

Payne, B. K., Jacoby, L. L., & Lambert, A. J. (2005). Attitudes as accessibility bias: Dissociating automatic and controlled components. In R. A. Hassin, J. S. Uleman, & J. A. Bargh, (Eds.), *The New Unconscious.* New York: Oxford University Press.

Payne, B. K., Lambert, A. J., & Jacoby, L. L. (2002). Best laid plans: Effects of goals on accessibility bias and cognitive control in race-based misperceptions of weapons. *Journal of Experimental Social Psychology, 38,* 384–396.

Richeson, J. A., Baird, A. A., Gordon, H. L., Heatherton, T. F., Wyland, C. L., Trawalter, S., & Shelton, J. N. (2003). An fMRI examination of the impact of interracial contact on executive function. *Nature Neuroscience, 6,* 1323–1328.

Richeson, J. A., & Shelton, J. N. (2003). When prejudice does not pay: Effects of interracial contact on executive function. *Psychological Science, 14,* 287–290.

Snodgrass, J. G., & Corwin, J. (1988). Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *Journal of Experimental Psychology: General, 117,* 34–50.

Srull, T. K., & Wyer, R. S. (1980). Category accessibility and social perception: Some implications for the study of person memory and interpersonal judgments. *Journal of Personality and Social Psychology, 38,* 841–856.

Tanner, W. P., & Swets, J. A. (1954). A decision making theory of visual detection. *Psychological Review, 61,* 401–409.

von Hippel, W., & Gonsalkorale, K. (2005). "That is bloody revolting!" Inhibitory control of thoughts better left unsaid. *Psychological Science, 16,* 497–500.

von Hippel W., Silver, L. A., & Lynch, M. E. (2000). Stereotyping against your will: The role of inhibitory ability in stereotyping and prejudice among the elderly. *Personality and Social Psychology Bulletin, 26,* 523–532.

Wittenbrink, B., Judd, C., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology, 72,* 262–274.

Received August 20, 2004
Revision received April 6, 2005
Accepted April 22, 2005 ■