Bringing Automatic Stereotyping Under Control: Implementation Intentions as Efficient Means of Thought Control

Brandon D. Stewart and B. Keith Payne

Pers Soc Psychol Bull 2008; 34; 1332 originally published online Jul 25, 2008;
DOI: 10.1177/0146167208321269

Bringing Automatic Stereotyping Under Control: Implementation Intentions as Efficient Means of Thought Control

Brandon D. Stewart
University of Queensland

B. Keith Payne
University of North Carolina, Chapel Hill

The evidence for whether intentional control strategies can reduce automatic stereotyping is mixed. Therefore, the authors tested the utility of implementation intentions—specific plans linking a behavioral opportunity to a specific response—in reducing automatic bias. In three experiments, automatic stereotyping was reduced when participants made an intention to think specific counterstereotypical thoughts whenever they encountered a Black individual. The authors used two implicit tasks and process dissociation analysis, which allowed them to separate contributions of automatic and controlled thinking to task performance. Of importance, the reduction in stereotyping was driven by a change in automatic stereotyping and not controlled thinking. This benefit was acquired with little practice and generalized to novel faces. Thus, implementation intentions may be an effective and efficient means for controlling automatic aspects of thought.

Keywords: *automatic stereotyping; control; conscious intentions; process dissociation; goals*

People sometimes do one thing when they intend to do another. This inconsistency in thought and behavior usually occurs because aligning actions with intentions requires ability and motivation, which are often in short supply. Distractions, fatigue, or other demands may limit people's ability and motivation and thus their control. In the absence of such intentional control, it is an important fact that we do not act randomly. We rely instead on well-learned habits, categories, or highly accessible thoughts (Bargh, 1999; Ouellette & Wood, 1998). These well-worn habits of thought and action are often useful and adaptive. Sometimes, however, these accessible habits can interfere with our goals.

In such cases, avoiding or overcoming automatic thinking is an essential component of achieving goals and avoiding unwanted biases. The influence of conscious goals in moderating automatic stereotyping has become an important focus of research (e.g., Bargh, 1999). A primary theoretical issue at stake is whether psychological effects deemed to be "automatic" also should be considered uncontrollable. Research to date has provided mixed evidence. Here, we briefly review some key points in this debate and then present new data bearing on whether, and under what circumstances, people can moderate their own automatic stereotype biases. We tested a special type of conscious control strategy, an implementation intention, to see whether it reduced automatic stereotyping.

Authors' Note: This research was conducted as part of the first author's doctoral dissertation under the supervision of Russ Fazio and Keith Payne. We would like to thank Russ Fazio and Kathleen Pierce for helpful comments on drafts of this manuscript and members of the Social Cognition Research Group at The Ohio State University for feedback during the development of this research. Please direct correspondence to Brandon D. Stewart, School of Psychology, University of Queensland, St. Lucia, 4072, Australia; e-mail: stewart.507@osu.edu.

PSPB, Vol. 34 No. 10, October 2008 1332-1345
DOI: 10.1177/0146167208321269

Malleability and Stereotyping

There is little doubt that people can control their behaviors when they have motivation and opportunity to think carefully (for review, see Fazio & Towles-Schwen, 1999). The question of interest is instead about more subtle, implicit forms of stereotyping that are likely to take place when it is difficult to control responses. Early research on implicit attitudes and stereotyping led researchers to conclude that automatic biases were likely to be very rigid and require arduous learning processes to change. More recent research has found that implicit measures of stereotyping may be quite malleable in response to contextual factors (for review, see Blair, 2002; Dasgupta & Greenwald, 2001; Kawakami, Dovidio, Moll, Hermsen, & Russin, 2000; Wittenbrink, Judd, & Park, 2001). In particular, performance on implicit measures has been shown to be moderated by a few factors, including (a) the context surrounding the stimulus (Barden, Maddux, Petty, & Brewer, 2004; Lowery, Hardin, & Sinclair, 2001; Wittenbrink et al., 2001), (b) situationally activated self-enhancement motives (Sinclair & Kunda, 1999; Spencer, Fein, Wolfe, Fong, & Dunn, 1998), (c) focus of attention (Macrae, Bodenhausen, Milne, Thorn, & Castelli, 1997), (d) effortful practice (Kawakami et al., 2000), and (e) promotion of counterstereotypes (Blair, Ma, & Lenton, 2001). There is little doubt that implicit measures of bias can be moved. It is much less clear, however, whether one's own intentions can be the primary mover.

Malleability and Intentional Control

Studies showing malleable responses on implicit tests have raised optimism that stereotyping might be more amenable to volitional control than previously thought. After all, if implicit measures are sensitive to mental images and varying situations, shouldn't they be responsive to participants' efforts to overcome bias? It is probably too soon, however, to equate situational malleability with controllability. In fact, there is mixed evidence on whether people's intentions, rather than the passive influence of context, can moderate automatic stereotyping. Some of the experiments that have manipulated people's awareness of stereotyping and active attempts at controlling automatic stereotyping have tended to show "backfire" effects, whereas others have found positive effects. One study, for example, found a decrease in race stereotyping measured with the Implicit Association Test when participants had a goal to avoid bias and the experimenter was both likable and African American (Lowery et al., 2001). This study suggested that the status or liking of the African American partner must be high for this effect to be observed.

Several other studies, however, have shown increased stereotyping on later tasks, when participants have the conscious goal to change their responses (Galinsky & Moskowitz, 2000; Liberman & Foerster, 2000; Macrae, Bodenhausen, Milne, & Jetten, 1994; Payne, Lambert, & Jacoby, 2002). In one study, participants were directly warned to avoid stereotype bias while performing an implicit task. The warned participants showed increased bias (Payne et al., 2002). Similar counterintentional effects have been shown when the motivation to control stereotypes was more subtle, such as the pressure imposed by actual social interactions (Lambert et al., 2003; Richeson et al., 2003; Richeson & Shelton, 2003). At least in some situations, it appears that attempts to control automatic stereotyping may actually set people on a path toward stereotyping, especially under conditions where control is difficult to achieve.

Because controlling automatic responses can be difficult, it is important to identify effective strategies for avoiding bias that people can employ on demand. It is especially beneficial to identify strategies that can be used with little effort given the added difficulty of meeting all conditions for correcting bias once it has occurred (i.e., awareness of bias, its direction and magnitude, and the motivation and ability to correct; Wilson & Brekke, 1994). Although some studies have demonstrated that conscious strategies can moderate automatic stereotyping, these strategies are not without limitations.

One study showed that implicit task performance could be altered by spending 5 min imagining a counterstereotypical person (e.g., strong woman; Blair et al., 2001). Participants who completed the counterstereotypical mental imagery showed less automatic stereotyping across a variety of implicit measures. Although this strategy reduced stereotyping, the mechanisms for this effect remain unclear. Participants may not have been aware of trying to avoid bias or control stereotyping, which often ironically induces people to stereotype. The question of whether malleability provides evidence for controllability remains open.

Another study that altered implicit test performance is especially relevant to the present studies. Blair and Banaji (1996) had participants complete a task in which they classified names as male or female when those names were preceded by stereotypical primes. These primes were either consistent or not consistent with the gender stereotype (e.g., aggressive—Mike; flowers—Tom). Participants showed less stereotypical associations when they had been informed to expect the stereotype-inconsistent combination. Although this finding is encouraging, it leaves several questions. First, this effect was driven mostly by the large increase in stereotyping when participants had stereotype-consistent

expectations. Second, the design of these experiments altered the contingencies between prime and target pairs. There was a greater percentage of counterstereotypical pairings in the stereotype-inconsistent expectation condition than in the stereotypical expectation condition. Participants may have learned this covariation and therefore been quicker in responding to counterstereotypical pairs within the experimental context (Barker & Andrade, 2006). The reduction in stereotyping could thus have been due to the disproportionately greater number of stereotype-inconsistent pairs rather than to intentions.

PRESENT RESEARCH

Efficient Goal Pursuit With Implementation Intentions

The present research seeks to better understand when strategic attempts at controlling bias will be successful, and what cognitive processes are required. Past research suggests that a certain strategy, an *implementation intention,* may be promising because it may allow people to "automate" goal pursuit (Brandstatter, Lengfelder, & Gollwitzer, 2001). Implementation intentions are if-then action plans (e.g., "When I leave work, I will go exercise at the gym") that have been shown to help people enact their goals more efficiently than general intentions (e.g., "I will exercise more," Brandstatter et al., 2001; Gollwitzer & Brandstatter, 1997). Implementation intentions are believed to function by increasing the strength of association between the environmental cue (i.e., the "if" portion) and the goal-directed intention, and by making both the environmental cue and the goal-directed intention more accessible (Aarts, Dijksterhuis, & Midden, 1999; Webb & Sheeran, 2007). Once a person has made the initial if-then intentions, these intentions require little motivation and effort to employ (Brandstatter et al., 2001). In contrast, general intentions are proposed to make the intention accessible, but not in response to the relevant environmental stimulus. Thus, general intentions may not produce the intended response efficiently.

Some research has begun to show that implementation intentions may reduce automatic stereotyping on reaction time measures, at least under specific conditions (Gollwitzer, Fujita, Oettingen, 2004). In this research, Gollwitzer and colleagues asked participants to complete a Stroop task in which two male names, two female names, and control letter strings were presented as primes, followed by target words (gender-stereotypical and gender-neutral attributes) that were presented in four different colors. One group of participants was assigned a goal intention to judge the women in a nonstereotypical manner, and others were assigned simple intentions about how they would respond when they encountered a specific person (i.e., "And, whenever I see Ina, I will ignore her gender!"). Creating an intention to ignore gender in response to a specific individual reduced automatic activation of gender on the Stroop task (i.e., no delays in color naming on gender-stereotypical words).

Although this research is promising, its breadth and utility could be expanded by overcoming two limitations. First, the implementation intentions used were limited to one specific person from the category. Although it may be helpful to avoid stereotyping Ina, the social import of stereotypes is that they draw sweeping generalizations based on broad social categories. The most useful antistereotyping strategy would be one whose remedy is equally sweeping, and applies to the entire category.

Second, the implicit and explicit measures used in previous research are assumed to be process pure and to tap automatic and conscious processes, respectively. However, implicit and explicit tasks are not likely to reflect purely automatic or controlled processing (Jacoby, 1991; Payne, Jacoby, & Lambert, 2004). For example, if people are uncertain about how to respond to a question on a self-report measure, they may answer with whatever comes to mind first; thus, more automatic processes can influence responses on an explicit measure. Controlled processing can also influence performance on implicit measures. In a word-stem completion measure, for example, people may respond with whatever comes to mind first. However, if nothing comes to mind, they may intentionally search their memory for a word that fits the stem. As a result, both automatic and controlled thinking may influence the responses that are assessed. In other words, changes in implicit task performance do not necessarily reflect changes in underlying automatic processing. Critically, in the previous research showing that implementation intentions can modify performance on an implicit task, there remains uncertainty as to whether implementation intentions functioned by (a) reducing the automatic activation of stereotypes, by (b) increasing the amount of attention or concentration (i.e., controlled thinking) used during the tasks, or by both processes.

Separating Automatic and Controlled Influences

Research using process dissociation methods has shown that performance on implicit tasks is due to both automatic and controlled thinking (Hense, Penner, & Nelson, 1995; Jacoby, 1991; Payne, 2005). Furthermore, recent research has shown that participants with high levels of cognitive control were less influenced by automatically activated stereotypes, even on implicit measures (Payne, 2005; Payne et al., 2004). Other research has shown that nervousness-induced reductions in controlled thinking actually increased the influence of

stereotypes (Lambert et al., 2003; Richeson & Trawalter, 2005). Overall, studies demonstrating malleability on implicit measures could be explained by changes in automatic or controlled components of responses. It therefore becomes important to seek converging evidence from multiple approaches, including ones that do not rely on task dissociations. The process dissociation procedure (PDP) provides such an alternative approach to separating automatic and controlled processing (Jacoby, 1991).

The PDP estimates automatic and controlled contributions to a single behavior by pitting intentionally controlled thinking against automatic thinking within a single task. A process is considered automatic if it drives behavior whether it was intended or not. Intentional control is measured as the difference between performance when a person intends to respond a certain way and performance when the person intends not to respond in that way. To the extent that people can produce a particular response when they intend to, and not produce that response when they intend not to, they are exercising control.

Consider the sequence of events that unfolds very quickly when a person sees someone with an object and tries to determine whether the object is a threat. Within a few hundred milliseconds, the typical perceiver is reacting differently to a young Black man than to a young White man (Macrae & Bodenhausen, 2000). Thoughts and feelings associated with these social categories are activated, potentially leading to biased responses. At the same time, the perceiver is also focusing attention on the object, trying to discern what it is, and trying to respond appropriately. One stream of information—stereotypical thoughts and feelings—could lead to an assumption that the object was dangerous regardless of what it truly is. The second stream of information—selective attention to the features of the object—would lead to correctly responding regardless of who the person is. The perceiver's ultimate behavior depends on the quick resolution of these different responses.

Consider now how a perceiver who has made a specific plan for what to think beforehand might respond. According to research on implementation intentions, people can carry out goals efficiently and with little effort by identifying a specific cue and a specific goal-appropriate action in an if-then format. In the present example, stereotypic thoughts and feelings of threat, criminality, and danger provide the link between race and weapons. To counter this stereotype, we might have the perceiver think its opposite, "safe." Having previously made this if-then plan, the chain of events might change drastically. Upon categorizing the person as a Black man, thoughts of safety would be activated instead of threat. The overall task goal to correctly identify weapons remains in place. But the implementation intention aims to change what thoughts are automatically activated.

In three experiments, we used implicit measures and process dissociation analyses to test whether counterstereotypical implementation intentions reduced stereotyping and whether the reduction occurred primarily through changes in automatic stereotyping, controlled responding, or a combination of the two. If making a counterstereotypical implementation intention altered the goal of identifying weapons or the attention to the features of the objects, then it should affect the controlled component. In contrast, if the intention altered the thoughts that potentially bias responses, then it should affect the automatic component. Our designs allowed us to evaluate both possibilities.

EXPERIMENT 1

In Experiment 1, we used a weapons-identification paradigm because it greatly restricts control and because process dissociation can be used to estimate the influences of automatic and controlled responding (Payne, 2001; Payne et al., 2002). The paradigm uses multiple criteria—including a short stimulus-onset asynchrony, speeded responses, and process dissociation analysis—for assessing automaticity. Previous research has validated that the automatic stereotyping estimate in this task can be characterized as an estimate of automatic processing because it (a) occurs rapidly and (b) occurs despite intentions to the contrary (Payne et al., 2002). Cognitive control may be conceptualized here as the ability to respond as intended (i.e., accurately identify the object based on the objective features of the stimulus and the task goals) without bias from stereotypes. The control estimate can be characterized as controlled thinking because it (a) is defined by the ability to respond in accordance with explicit goals and (b) is sensitive to the manipulation of processing time.

We tested the effectiveness of an implementation intention to think "safe" in response to Black faces. To insure that any differences between conditions were not due only to having any kind of implementation intention, we compared this plan against matched intentions to either think "quick" or think "accurate." We chose these comparison intentions because the task already requires participants to respond quickly and accurately. Therefore, these thoughts should not differ from the thoughts already activated by the task. In contrast, Black faces typically activate thoughts of threat. The think-"safe" intention should change that thought and hence change responses more than the other two intentions.

Method

Participants

One hundred forty-six participants (63 women and 83 men) from introductory psychology courses participated for course credit.

Design and Procedure

Weapons task. The experiment used a 3 (implementation intention: think safe, think quick, think accurate) × 2 (race prime: African American, Caucasian face) × 2 (object: handgun, hand tool) design with race prime and object as within-participant factors. Participants were informed that the experiment concerned how people make simple but quick decisions. They were told that they would see pairs of pictures flashed one after the other. The first picture would be a face. Initially, participants were told that they should do nothing with the face because it was simply a warning signal that the second image was about to appear. The second image would be an object, and their job was to identify the object as quickly and accurately as possible; thus, all participants had a general goal to respond quickly and accurately. A second screen then appeared, informing them that they should judge what each object was by pressing the "gun" or "tool" key on the keyboard. Participants completed 32 practice trials to become acquainted with the task, before the critical trials began. On each trial, a message appeared that asked them to respond more quickly if they responded too slowly.

The primes included black-and-white photographs of four African American male and four Caucasian male faces. The target object was either a handgun or a hand tool. The prime remained on the screen for 200 ms, and it then was replaced by the object for 100 ms. Thus, the stimulus onset asynchrony was 200 ms. After the object was presented, it was replaced with a visual mask that remained on the screen until the participant responded. As in previous research, participants were warned to respond more quickly if their response was longer than 450 ms (Payne et al., 2002). For each trial, the prime appeared 250 ms after the previous response.

Intentions manipulation. Following the basic instructions and practice trials, participants received the implementation intention manipulation. All participants were told that the race of the face affects the way that people classify the second object. They then were informed that

> Studies on perceptual readiness have shown that a certain mental exercise helps increase one's reactivity. If you make a resolution to respond in a particular way to a specific type of face, you will be able to change your reaction.

Participants were shown four cards on the computer monitor and were asked to "Choose the type of face you will respond to by pressing the 1, 2, 3, or 4 key." We used this card-selection procedure because previous research has shown that giving participants a feeling of choice increases commitment to the goal. All participants were told that they had selected to respond to African American faces.

The instructions for each condition diverged at this point. Participants in the counterstereotypical condition were told,

> We would like you to commit yourself to responding to the Black face by thinking the word "safe." By thinking the word "safe," you are reminding yourself on each trial that you are just as safe interacting with a Black individual as with a White individual.

Participants were then instructed,

> In order to firmly commit yourself to responding to the Black face, please say to yourself silently, "I definitely want to respond as accurately as possible by thinking the word, safe. Whenever I see a Black face on the screen, I will think the word, safe."

Past research has successfully used an identical format to provide participants with implementation intentions (Brandstatter et al., 2001).

Participants in the "accurate" condition were told,

> We would like you to commit yourself to responding to the Black face by thinking the word "accurate." By thinking the word "accurate," you are reminding yourself on each trial that you should accurately identify the objects that appear after the Black and White faces.

They were then instructed,

> In order to firmly commit yourself to responding to the Black face, please say to yourself silently, "I definitely want to respond as accurately as possible by thinking the word, accurate. Whenever I see a Black face on the screen, I will think the word, accurate."

Participants in the final condition were asked to think the word "quick" whenever they saw a Black face. Both the think-"quick" and the think-"accurate" implementation intentions were task relevant but were not associated with tools or weapons. We therefore expected that they would not interfere with the task of classifying objects quickly and accurately. Furthermore, all participants were asked to respond quickly and accurately as part of the task instructions. Finally, participants completed 192 critical trials, divided into three blocks of 64 trials.

The instructions used in this experiment replicate the if-then format that is theorized to be the driving force behind the effectiveness of implementation intentions. According to the theory, an if-then format links an environmental stimulus with a specific response (e.g., when I see a Black face, I will then think "safe"). The creation of this link increases the accessibility of both the environmental stimulus and the specific thought or intention to which it was linked. Although we used "when-then" phrasing rather than "if-then," it would make little sense to participants to use "if" rather than "when" in the context of a task in which it was clear that they would see the faces repeatedly. Gollwitzer and colleagues have successfully used this same wording in previous studies (e.g., Brandstatter et al., 2001).

Results

The dependent measure of interest was error rates: the misidentification of tools as guns and of guns as tools. We performed a 3 (implementation intention) × 2 (race prime) × 2 (object) repeated-measures ANOVA, with prime and object as repeated measures. We observed the expected Race Prime × Object interaction in which participants falsely identified a tool as a gun more often after being primed with a Black compared to a White face, and falsely identified a gun as a tool more often after a White face, $F(1, 143) = 55.91$, $p < .001$. This pattern represents stereotypic responding. We also observed the predicted three-way interaction, indicating that the degree of stereotyping varied across implementation intention conditions, $F(1, 143) = 3.86$, $p = .02$. The means are presented in Figure 1.

The three-way interaction was driven by the fact that the two-way Race Prime × Object interaction was significantly weaker in the safe intention group, $F(1, 49) = 3.83$, $p = .06$, $\eta^2 = .07$, than in the quick intention group, $F(1, 49) = 33.60$, $p < .001$, $\eta^2 = .41$, or the accurate intention group, $F(1, 45) = 31.88$, $p < .001$, $\eta^2 = .42$. When only the quick and accurate intention groups were analyzed, these groups did not differ from each other; the type of implementation intention did not moderate the effects, $F(1, 94) = 0.69$, $p = .41$, $\eta^2 = .01$. These analyses show that the significant effect of implementation intentions was driven by decreased stereotypic responding in the safe intention condition.

Simple effects analyses showed that in the quick intention group, there were more false "gun" responses after a Black prime than after a White prime, $F(1, 49) = 9.74$, $p = .003$. The same effect was found in the accurate intention group, $F(1, 45) = 20.56$, $p < .001$. In the safe intention group, however, there was no significant difference in false "gun" responses after Black versus White primes, $F(1, 49) = 1.61$, $p = .21$.

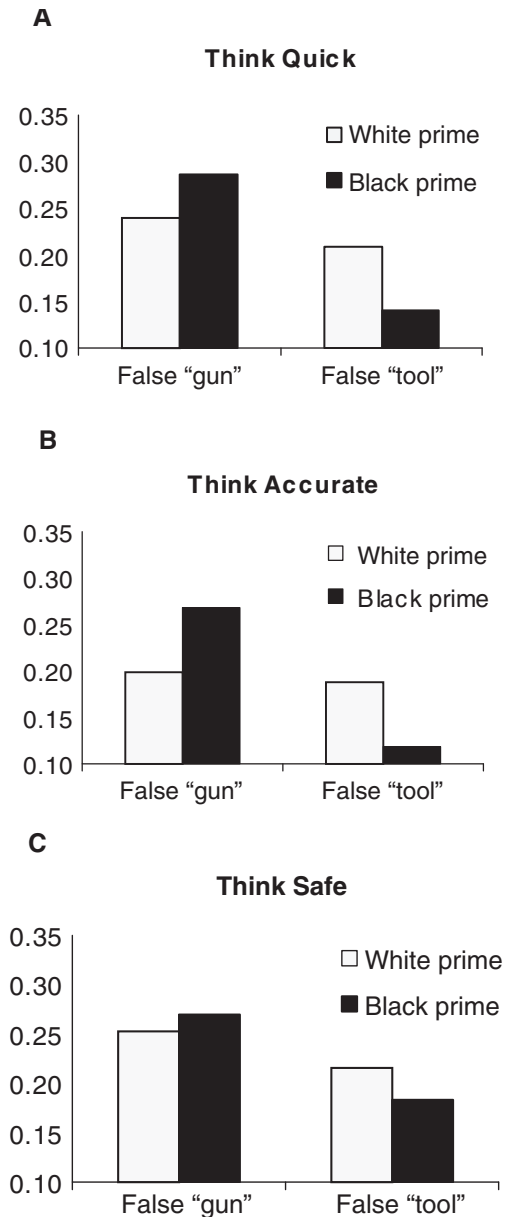

Figure 1 Proportion of Errors by Type of Object Error (False "Gun" vs. False "Tool" Response), Race Prime (White vs. Black), and Implementation Intention Condition in Experiment 1.

The tendency to mistakenly call a gun a tool was greater after White primes than Black primes in all three groups. However, that effect was stronger in the quick intention group, $F(1, 49) = 35.05$, $p < .001$, and the accurate intention group, $F(1, 45) = 26.64$, $p < .001$, than in the safe intention group, $F(1, 49) = 4.11$, $p = .05$. Together, these results support the hypothesis that the counterstereotypical implementation intention to think "safe" effectively reduced the impact of stereotypes.

Process Dissociation Analysis

We used the PDP to separate automatic bias (A) and controlled responding (C), as was done in Payne (2001). The object-identification task included both *congruent* conditions, in which automatic stereotyping and controlled thinking act in the same direction, and *incongruent* conditions, in which they oppose one another. In a congruent condition, responding either on the basis of the actual object or responding on the basis of the racial category, in the absence of control, will lead to a correct response. This relationship may be expressed mathematically in the following equation: P(correct | congruent) = C + A(1 – C). Therefore, the probability of responding correctly on congruent trials is the probability of controlled thinking, C, plus the probability of an automatic association when control fails, A(1 – C). In the incongruent conditions, automatic influences and controlled thinking should lead to contradictory responses. Stereotypical errors will result to the extent that control fails, and participants respond based on the activated stereotype. Mathematically, this can be written as the P(stereotypic error | incongruent) = A(1 – C). The two sets of equations then may be used to derive estimates of each process.

Control is solved algebraically as the difference between correct responses in the congruent condition and errors in the incongruent condition: C = P(correct | congruent) – P(stereotypic error | incongruent). Given the estimate of control, one can solve for the automatic stereotype bias estimate. This automatic bias is equal to the probability of stereotypical errors in the incongruent condition divided by failures to control: A = P(stereotypic error | incongruent) / (1 – C). The automatic estimate was scored such that higher numbers represent a greater bias toward responding "gun." Automatic stereotyping is observed when automatic bias is higher after a Black prime than after a White prime.

Automatic estimates were analyzed using a 2 (race prime) × 3 (intention condition) repeated-measures ANOVA. The means are reported in Table 1. There was a significant main effect of race prime, $F(1, 143) = 76.69$, $p < .001$, $\eta^2 = .35$, showing that the automatic bias toward responding "gun" was higher following Black primes than White primes. This main effect was qualified by the predicted interaction, $F(2, 143) = 8.51$, $p < .001$, $\eta^2 = .11$. Simple effects tests showed that race primes significantly affected the automatic estimates in the quick intention condition, $F(1, 49) = 48.24$, $p < .001$, and the accurate intention condition, $F(1, 45) = 44.12$, $p < .001$, but not in the safe intention condition, $F(1, 49) = 2.88$, $p = .10$. Thus, the automatic component of stereotypical responding was reduced to nonsignificance by the think-"safe" implementation intention.

TABLE 1: Automatic Estimates by Prime and Implementation Intention (Experiment 1)

	Race Prime			
	Black		*White*	
Implementation Intention Condition	M	SD	M	SD
Safe	$.58_a$	.17	$.53_a$	.18
Quick	$.68_b$	.16	$.51_a$	.17
Accurate	$.68_b$	.16	$.48_a$	.17

NOTE: Estimates are probabilities (theoretical range from 0.0 to 1.0) such that values represent the probability of making a "gun" response when control fails. Means with different subscripts within a row differ at $p < .05$.

TABLE 2: Cognitive Control Estimates by Prime and Implementation Intention (Experiment 1)

	Race Prime			
	Black		*White*	
Implementation Intention Condition	M	SD	M	SD
Safe	$.55_a$	.28	$.53_a$	.30
Quick	$.57_a$	.21	$.55_a$	.23
Accurate	$.61_a$	.23	$.61_a$	.21

NOTE: Estimates of control can range between a theoretical minimum and maximum of 0.0 and 1.0, with 1.0 representing perfect control in identifying guns and tools and 0.0 representing chance responding.

Estimates of control were analyzed using the same ANOVA model. The means are shown in Table 2. There were no significant main effects or interactions, showing that neither race primes nor implementation intentions influenced the controlled component of responses (all ps ≥ .20). The control estimate is based on the ability to discriminate between gun and tool targets. The lack of effects on control suggest that the different intentions did not alter participants' overall task goals to distinguish guns and tools. Their control over responses was not affected by the intention to think "safe." Instead, that thought altered how they responded when controlled responding failed.

One might speculate that the think-"safe" implementation intention reduced automatic bias by causing participants to simply slow down and respond more cautiously when they were presented with a Black face. This explanation, however, was not supported by the data. An ANOVA on log transformed reaction times showed that participants in the counterstereotypical condition did not slow down after Black faces ($M = 335$ ms) relative to White faces ($M = 340$ ms), $F < 1$.

Moreover, an ANOVA with object (gun or tool) as a within-participant factor and intention as a between-participant factor was performed on reaction times for Black prime trials; this analysis showed that participants in the counterstereotypical condition (M = 335 ms) did not slow down after Black faces when compared to quick intention participants (M = 333 ms), $F < 1$. Nor was the Intention × Object interaction significant, $F < 1$. Thus, participants with a counterstereotypical intention did not use a response strategy in which they took longer to respond after Black primes.

Discussion

This experiment provides support for the idea that implementation intentions reduce stereotyping by reducing automatic effects of the stereotype. The process dissociation results demonstrated that the counterstereotypic implementation intentions reduced automatic bias but not controlled processing (i.e., responding as intended). Moreover, they did so rapidly and without much effort or increased caution on the part of the participants, as demonstrated by the reaction-time analyses.

The finding that only automatic bias was reduced is particularly striking because many experiments using the weapon paradigm have demonstrated that controlled processing can be increased (e.g., by taking longer), but none has shown that the automatic stereotyping bias can be reduced (Lambert et al., 2003; Payne, 2001; Payne et al., 2002). Our research is among the first to support the idea that a counterstereotypical implementation intention can be an effective strategy for intentionally avoiding automatic stereotyping. In contrast to previous findings, which showed that deliberate attempts to avoid prejudice often increase stereotyping (Monteith, Ashburn-Nardo, Voils, & Czopp, 2002; Payne et al., 2002; Richeson & Trawalter, 2005), we found evidence that implementation intentions might be effective while avoiding this problem.

Although these results demonstrate that counterstereotypical implementation intentions are an effective means for reducing automatic stereotyping, the scope of the effect is not yet known. On one hand, implementation intentions might affect people's reactions only to the specific individuals about whom the intention is made. On the other hand, the effect might be more general. The implementation intentions used here identified the targets at the level of an entire social group, rather than the specific individuals used to represent that group. Therefore, the implementation intentions may extend to other African American individuals as well. From the perspective of practical applications, a broader effect would be much more useful because stereotyping itself operates at the group level. To examine the scope and generality of this technique, in the next experiment we tested how well the effect generalized to a new set of faces. Furthermore, we reduced the complexity of the implementation intention instructions in order to demonstrate that a brief implementation intention instruction could be effective. Finally, because the think-"quick" and think-"accurate" conditions did not differ in Experiment 1, we included only the think-"quick" condition in Experiment 2 as a comparison for the think-"safe" intention.

EXPERIMENT 2

In Experiment 2, we again used the weapon-identification paradigm. However, we added a manipulation in which previously unseen faces were presented to participants, and we tested how well the counterstereotypical implementation intentions generalized (e.g., reduced stereotyping in response to the new faces). The first phase of the experiment was the same as in Experiment 1, but in the second phase, the priming faces were subtly switched to a new set of Black and White faces. If the effects we have observed are specific to the set of exemplars on which participants have practice using implementation intentions, then we can expect implementation intentions to reduce stereotyping for the original faces but not for the new faces. In contrast, if the implementation intentions influence responses to the whole social category, then we can expect that they will reduce stereotyping for both original and new exemplars.

Method

Participants

One hundred twenty-five participants (63 women, 62 men) participated for course requirement.

Design and Procedure

This experiment used a similar manipulation of implementation intentions, the same weapon-identification paradigm, and the same instructions as Experiment 1, with the following exceptions. First, participants were randomly assigned to a think-"quick" or think-"safe" intention group. The implementation intention instructions were simplified to "In order to firmly commit yourself to responding to the Black face by thinking 'safe,' please say to yourself silently, 'I definitely want to respond to the Black face by thinking safe.'" Parallel wording was used for the think "quick" group. Second, the 192 critical trials were divided into four blocks of

48 trials instead of three blocks of 64 trials. Participants committed to their implementation intention, as was done in the previous experiment, and then proceeded to complete 16 practice trials and one block of critical trials using the same four African American and four Caucasian faces as primes.

Following this first critical block, participants completed a second and a third critical block with novel-face primes that had not been used in the previous blocks. Following these generalization blocks, participants completed a final block of trials that used the original faces from the first block of trials. The final block was included to control for the effects of practice. Other research has found that participants may show decreases in stereotyping in weapon identification with practice (Plant, Peruche, & Butz, 2005). If the generalization block was the final block, then any reduction in stereotyping could be explained by practice effects. Such practice effects would have a greater influence on the fourth block than on the second and third blocks. Thus, a design with a final nongeneralization block allows the influence of practice to be distinguished from the generalization to new faces. The experiment used a 2 (implementation intention: think safe, think quick) × 2 (race prime: African American, Caucasian face) × 2 (object: handgun, hand tool) × 4 (block) design with prime, object, and block as within-participant factors.

Results and Discussion

As a first analysis, we tested the influence of implementation intentions on performance, ignoring the factor of novel versus practiced faces, in order to establish whether the findings from Experiment 1 replicated. Error rates were analyzed using a 2 (implementation intention: safe, quick) × 2 (race prime: African American, Caucasian) × 2 (object: gun, tool) ANOVA, with race prime and object as repeated measures. We observed the expected Race Prime × Object stereotyping interaction in which participants falsely identified a tool as a gun more often after being primed with a Black compared to a White face, and falsely identified a gun as a tool more often after a White face, $F(1, 123) = 35.41$, $p < .001$, $\eta^2 = .22$.

Most important, replicating Experiment 1, this interaction was qualified by the predicted three-way Race Prime × Object × Implementation Intention interaction, $F(1, 123) = 12.16$, $p = .001$, $\eta^2 = .09$, as displayed in Figure 2. Participants in the quick condition showed significant stereotyping, as indicated by the Race Prime × Object interaction, $F(1, 61) = 38.97$, $p < .001$, $\eta^2 = .39$. In contrast, those in the safe implementation intention condition showed much less evidence of stereotyping, as the Race Prime × Object interaction was smaller and only marginally significant, $F(1, 62) = 3.52$, $p = .07$, $\eta^2 = .05$.

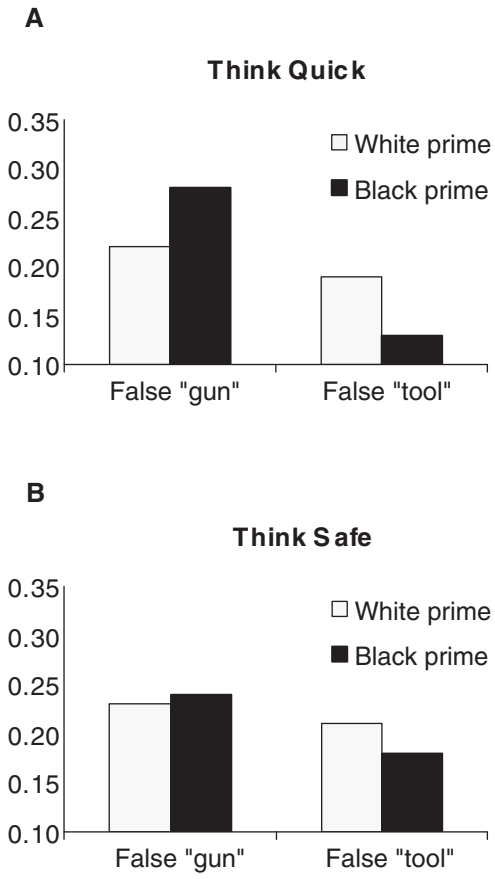

Figure 2 Proportion of Errors by Type of Object Error (False Gun vs. False Tool Response), Race Prime (White vs. Black), and Implementation Intention Condition in Experiment 2.

The analysis was repeated next with the block variable included to test whether the effectiveness of the implementation intentions differed across blocks. Block did not qualify the results, as evidenced by the lack of a four-way interaction, $F(3, 121) = 0.60$, $p = .62$, $\eta^2 = .02$. To more clearly graph the effect of different blocks on performance, the Race Prime × Object interaction was converted into a single performance bias score. This was the proportion of stereotype-consistent errors (false gun on Black prime trials + false tool responses on White prime trials) minus the proportion of stereotype-inconsistent errors (false tool responses on Black prime trials + false gun on White prime trials). Higher values on this score represent greater stereotyping. As shown in Figure 3, the effect of the counterstereotypical implementation intention was uniform across blocks. Of importance, this result indicates that the counterstereotypical intention generalized to the new faces inserted into Blocks 2 and 3.

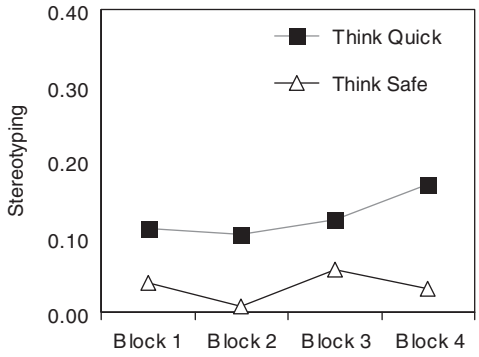

Figure 3 Stereotypical Bias as a Function of Block and Implementation Intention Condition in Experiment 2.

TABLE 3: Automatic Estimates by Prime and Implementation Intention (Experiment 2)

	Race Prime			
	Black		*White*	
Implementation Intention Condition	M	SD	M	SD
Safe	$.58_c$	.19	$.51_a$	.17
Quick	$.70_b$	.15	$.51_a$	.17

Automatic and Controlled Estimates

Automatic estimates were computed so that higher numbers represented a greater bias toward responding "gun." Stereotyping is observed when this gun bias is higher after a Black prime than after a White prime. A repeated-measures ANOVA on automatic bias, with race as a within-participant factor and implementation intention as a between-participant factor, revealed a significant Race Prime × Implementation Intention interaction, $F(1, 123) = 9.29$, $p = .003$. As shown in Table 3, simple effects tests showed that automatic stereotype bias was higher for Black primes than White primes in the quick intention condition, $F(1, 61) = 42.29$, $p < .001$, $\eta^2 = .41$, and in the safe condition, $F(1, 62) = 7.28$, $p < .05$, $\eta^2 = .11$; however, the difference was significantly smaller in the safe intention condition.

The effect of implementation intentions on automatic bias was not qualified by block, confirming that the reduction in automatic bias was uniform across novel and practiced exemplars, $F(3, 121) = 1.21$, $p = .31$, $\eta^2 = .03$. Furthermore, no differences between safe and quick conditions were observed for the estimate of controlled responding, $F(1, 123) = 0.63$, $p = .43$, or on self-reported amount of effort on the task, $F(1, 123) = 0.93$, $p = .34$. These automatic and controlled process results were consistent with results from the previous study. The counterstereotypical implementation intention to think "safe" reduced automatic stereotyping and did so quickly (i.e., in the first block of trials).

EXPERIMENT 3

Experiments 1 and 2 provided evidence that counterstereotypical intentions reduced automatic stereotyping within the weapon-identification task. In Experiment 3, we sought convergent evidence by using a response time measure, the Implicit Association Test (IAT). Replicating the effect of implementation intentions with a different validated implicit measure would demonstrate the robustness of the effect (Nosek, Greenwald, & Banaji, 2005).

Method

Participants

One hundred nineteen participants (54 women, 65 men) completed the study for course credit.

Design and Procedure

Participants were informed that they would be performing a categorization task in which they would be classifying words and images into groups, and that they should try to classify the items quickly. As is standard procedure for the IAT, participants categorized four types of stimuli (i.e., photos of Black and White male faces and pleasant and unpleasant words) using two response keys. The race stimuli included 135 × 140 mm black-and-white photographs of five Black and five White faces. The pleasant and unpleasant words were *freedom*, *peace*, *love*, *paradise*, *caress*, *sickness*, *death*, *filth*, *terrible*, and *poison*. These words were selected from stimuli used in previous IAT research, and the positive and negative words were equated for word length and extremity (Greenwald, McGhee, & Schwartz, 1998). Participants completed 20 practice trials of categorizing the faces as either Black or White, 20 practice trials categorizing the words as pleasant or unpleasant, and 20 practice trials categorizing words and images into all four categories.

Participants received the implementation intention manipulation following these practice trials but before the critical trials. The implementation intention manipulation was similar to Experiment 2 but was tailored to be appropriate to the task. Because the IAT measures associations between concepts and good versus bad evaluations, participants in the counterstereotypical implementation intention condition were told, "In order to firmly commit yourself to responding to the Black face by thinking 'good,' please say to yourself

silently, I definitely want to respond to the Black face by thinking 'good.'" Participants in the control condition were given the same instruction except with a goal to think "quickly."

Following the implementation intention manipulation, participants completed 40 critical trials of the IAT in which White and pleasant were associated with one response key and Black and unpleasant were associated together on another key; this block of 40 trials represented the stereotype-congruent data collection block. Next, the pairing of race and evaluative words was reversed, and after 40 practice trials participants completed 40 stereotype-incongruent critical trials. The order of these stereotype-congruent and incongruent blocks was counterbalanced. The design was thus a 2 (implementation intention: think good, think quickly) × 2 (stereotype compatibility: compatible, incompatible) design. Following the IAT, participants were asked to rate how much effort they expended on the task, from 1 = *did not try at all* to 5 = *tried extremely hard*.

Results

Data Reduction

The data were reduced based upon procedures used in previous research using the IAT D measure (Greenwald, Nosek, & Banaji, 2003). First, trials with latencies greater than 10,000 ms were removed from the data. Second, participants, with more than 10% of their response latencies less than 300 ms were removed. Third, error penalties were computed in which each error latency was replaced with the block mean plus 600 ms. Finally, the measure's scale was created by using individual-respondent standard deviations. One deviation from the scoring procedure was made. Because the manipulation of the independent variable occurred after the practice trials, only data from the critical congruent and incongruent blocks were used for computing the D score. The D measure was computed by subtracting the mean latency for the congruent block from the mean latency for the incongruent block and dividing this difference score by the pooled standard deviation (Greenwald et al., 2003). Higher values represent greater associations between Black and bad and between White and good.

D-Measure Analyses

A one-way ANOVA on this D measure revealed a main effect for type of implementation intention, $F(1, 118) = 9.05$, $p = .003$. As expected, the good intention significantly reduced stereotypical responses ($M = .084$, $SD = .276$) in comparison to the quick intention ($M = .246$, $SD = .312$).

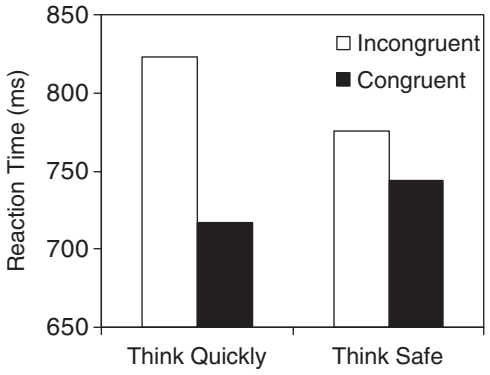

Figure 4 Reaction Time (ms) as a Function of the Stereotype Congruency of Race Prime and Word Pairings (Stereotype Congruent vs. Stereotype Incongruent), and Implementation Intention Condition.

To examine this effect more thoroughly, we conducted a repeated-measures ANOVA on reaction times (RTs) with congruency (RTs on congruent/incongruent blocks) as a within-participant factor and type of implementation intention as a between-participants factor. We observed a significant Congruency × Implementation Intention interaction, $F(1, 117) = 9.76$, $p = .002$, in which counterstereotypical implementation intentions produced a marginal speed-up on stereotype-incongruent trials, $F(1, 117) = 3.49$, $p = .06$, and a nonsignificant slowdown on stereotype-congruent trials, $F(1, 117) = 2.08$, $p = .15$ (see Figure 4). Overall, these results showed that the counterstereotypical intention to think "good" caused responses to be less stereotypical.

We next addressed the potential alternative explanation that decreased stereotyping occurred because the counterstereotypical intention motivated participants to respond more carefully. Such an explanation would suggest that participants might respond more slowly on the IAT overall, and would report greater effort. There was, however, no difference between conditions on the overall reaction time (quickly: $M = 770$, $SD = 106$; good: $M = 760$, $SD = 98$), $F(1, 118) = 0.32$, $p = .57$, nor on self-reported effort (quickly: $M = 3.54$, $SD = 1.08$; good: $M = 3.54$, $SD = 1.05$), $F(1, 118) < 0.001$, $p = .98$.

Discussion

Counterstereotypical implementation intentions decreased the influence of stereotypic associations on responses and seem to have done so relatively automatically. Both the reaction time data and the indirect assessment of associations provide converging evidence with the PDP studies, and they confirm the idea that the counterstereotypical implementation intentions operated

at an automatic level. One way that participants could have improved performance on the task would have been to slow their responses to all stimuli so as to respond more carefully. However, a general slowing of responses was not observed in the experimental condition. Thus, increased effort is an unlikely explanation for the reduction in stereotyping.

GENERAL DISCUSSION

The abundant evidence supporting the malleability of automatic stereotyping presents an intriguing question about the potential role of conscious goals in controlling bias. Is the "cognitive monster" of automatic stereotyping indeed chainable by conscious goals when opportunity to think carefully is limited? An early review of the automatic stereotyping literature indicated that there was not much evidence to support people's ability to control the influence of automatically activated stereotypes, and the few promising results required many, many repetitions to learn nonstereotypical responses (Bargh, 1999). Furthermore, the review argued persuasively that meeting all conditions for reducing bias (i.e., awareness of bias, awareness of the magnitude and direction, and motivation and ability to control bias) would occur simultaneously only on rare occasions. Overall, the promise of control did not look bright at that point in time. Shortly after this review, research began to show that automatic or implicit forms of stereotyping were influenced by manipulating the context in which the stimuli appeared, the configuration of the stimuli, or the extent of training to negate stereotypes (Barden et al., 2004; Blair et al., 2001; Kawakami et al., 2000; Lowery et al., 2001; Macrae et al., 1997; Spencer et al., 1998; Wittenbrink et al., 2001).

Although such evidence of situational malleability is suggestive of intentional control, there still is mixed evidence showing that people can intentionally control automatic stereotyping, especially in conditions that make control or monitoring responses difficult. For example, some experiments that indirectly induced participants to avoid stereotyping (e.g., interracial interaction; public situation) and those that gave participants an explicit goal to avoid bias have shown increased stereotyping (Lambert et al., 2003; Payne et al., 2002; Richeson et al., 2003; Richeson & Shelton, 2003), whereas a few others indicated that stereotyping may be reduced in certain social interactions (Lowery et al., 2001) or through extensive training (Kawakami et al., 2000). In the current research, we examined the influence of a conscious control strategy on automatic stereotyping by providing participants with a specific control strategy that required little effort and that they could employ on demand.

We demonstrated that stereotyping was reduced when participants used a counterstereotypical implementation intention. Of importance, this effect was driven by a change in automatic stereotyping. Cognitive control was not increased, as would be expected if participants directed more attention to or concentrated harder on the task. In Experiment 1, we used a weapon-identification task to greatly restrict participants' ability to control responses and to determine the separate influences of automatic and controlled thinking within the task. We found that counterstereotypical implementation intentions reduced stereotyping on this challenging task, and this outcome was due to a significant reduction in automatic stereotype bias as assessed by a process dissociation analysis. This finding is notable because it was the first to show a decrease in automatic bias on this task when participants were aware of bias (Lambert et al., 2003; Payne, 2001; Payne et al., 2002). Furthermore, the counterstereotypical implementation intention was more effective than an accuracy implementation intention. Next, we conducted a second experiment that ruled out the explanations of practice effects and that participants were responding to the specific faces they saw when forming their implementation intentions instead of responding to the category of African American. In Experiment 2, counterstereotypical implementation intentions generalized to Black faces that participants had not seen in previous trials or while forming their implementation intention.

A third experiment used an indirect assessment of stereotypic associations (i.e., IAT) in order to provide convergent evidence with the PDP measures. The results of this experiment confirmed that the counterstereotypical intention did reduce automatic stereotyping and that this reduction occurred because participants made less stereotypical responses; they slowed down on stereotype-incongruent trials and sped up on stereotype-congruent trials.

An important feature of the present line of research was its use of process dissociation, a technique that does not rely on task dissociations to estimate automaticity and control. It is important to use methods that do not rely solely on task dissociations to distinguish between automatic and controlled processing, because both types of processing can have separate influences within a given task and most implicit measures do not and cannot disentangle these influences (Jacoby, 1991; Payne, Jacoby, & Lambert, 2005). In regard to this problem of separate influences, research has found that variations in cognitive control do influence performance on implicit measures of stereotyping (Payne, 2005). This research found that low cognitive control allows automatic processing to have a greater impact on responses. Subjective feelings also have been shown to influence

implicit measures (Gawronski & Bodenhausen, 2005). Given these findings of the influence of controlled processing and subjective feelings on implicit tasks, the PDP provides important new insights by distinguishing automatic and controlled contributions.

Alternative Explanations and Potential Mechanisms

The results of Experiments 1 and 2 verified the hypothesis that providing participants with a counterstereotypical response strategy was effective because it decreased automatic bias. Across all three studies, participants were focused on the category of race and potential bias, but those with a counterstereotypical implementation intention were able to avoid automatic stereotyping. One may wonder if this strategy worked because participants were distracted from the relevant category, thus, the stereotype was never made accessible. This explanation is unlikely because the think-"quick" control intention should have been just as distracting as the counterstereotypical intention; therefore, it should have produced nonstereotypical responses similar to the counterstereotypical intention instead of causing more stereotyping. Moreover, Experiments 1 and 2 provided evidence that the social category of the faces was influencing participants in the counterstereotypical condition; thus, these participants were not distracted from the social category information.

The data from Experiments 1 and 2 also ruled out general priming of "safe" as an alternative explanation. An explanation based upon general semantic priming of "safe" in the counterstereotypical condition would predict that automatic bias toward responding "gun" would decrease after both Black and White primes in comparison to a control condition. But instead, bias was lower after Black primes, with no effect on White-prime trials. This observed decrease in the accessibility of the stereotype raises the interesting possibility that lateral inhibition or negative priming is the mechanism by which counterstereotypical implementation intentions function. Lateral inhibition refers to the unintentional disruption or inhibition of thoughts when those thoughts have been either ignored previously or when one focuses on different thoughts (Bodenhausen & Macrae, 1998; Tipper, 2001). In regard to the current research, focusing on a counterstereotypical thought may cause the relevant stereotypical thought (e.g., "dangerous" or "criminal" in the weapon paradigm) to be inhibited. Thus, it may be necessary for the specific intentions chosen to be connected to the task-relevant stereotype or social category. Participants with a quick or accuracy implementation intention experienced the same procedure as those with a counterstereotypical intention, but these control intentions did not reduce automatic stereotyping. Because the only difference between the conditions was the control intentions being unrelated to the task-relevant stereotype, it is possible that the thought must be connected to the stereotype in order for inhibition to occur. Future research will need to examine this lateral inhibition possibility more thoroughly.

Conclusion

Past research has found evidence supporting both the lack of controllability of stereotyping and the moderation of stereotyping. Our research clarifies one way in which conscious strategies can be used to overcome automatic stereotyping even when all the right circumstances (e.g., opportunity for controlled thinking, awareness, etc.) are not in place. It emphasizes that stereotyping due to difficult conditions may be surmounted if in advance of an interaction or situation one can create intentions that link a cognitive response to the relevant situation or group. To place these findings in context, the majority of factors known to moderate automatic stereotyping still involve variations in context or responses to context when participants are not alerted to potential bias. Perhaps the most effective approach would change both the social context (increasing positive intergroup contact, friendship, identity salience, etc.) and personal responses to that environment (as the implementation intentions did here). Of course, societal change is slow and beyond the control of any one individual, and the obstacles to changing large-scale social thinking are well known. Much less appreciated is the difficulty of changing one's own automatic thinking, which may be almost as difficult as changing environments and other people's thinking. Our studies suggest that implementation intentions are a good place to start.

REFERENCES

Aarts, H., Dijksterhuis, A. P., & Midden, C. (1999). To plan or not to plan? Goal achievement of interrupting the performance of mundane behaviors. *European Journal of Social Psychology*, *29*, 971-979.

Barden, J., Maddux, W. W., Petty, R. E., & Brewer, M. B. (2004). Contextual moderation of racial bias: The impact of social roles on controlled and automatically activated attitudes. *Journal of Personality & Social Psychology*, *87*, 5-22.

Barker, L. A., & Andrade, J. (2006). Hidden covariation detection produces faster, not slower, social judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*, 636-641.

Bargh, J. A. (1999). The cognitive monster: The case against the controllability of automatic stereotype effects. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 361-382). New York: Guilford.

Blair, I. V. (2002). The malleability of automatic stereotypes and prejudice. *Personality and Social Psychology Review*, *6*, 242-261.

Blair, I. V., & Banaji, M. R. (1996). Automatic and controlled processes in stereotype priming. *Journal of Personality and Social Psychology*, *70*, 1142-1163.

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

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

Brandstatter, V., Lengfelder, A., & Gollwitzer, P. M. (2001). Implementation intentions and efficient action initiation. *Journal of Personality and Social Psychology*, *81*, 946-960.

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

Fazio, R. H., & Towles-Schwen, T. (1999). The MODE model of attitude-behavior processes. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 97-116). New York: Guilford.

Galinsky, A. D., & Moskowitz, G. B. (2000). Perspective-taking: Decreasing stereotype expression, stereotype accessibility, and in-group favoritism. *Journal of Personality & Social Psychology*, *78*, 708-724.

Gawronski, B., & Bodenhausen, G. V. (2005). Accessibility effects on implicit social cognition: The role of knowledge activation and retrieval experiences. *Journal of Personality and Social Psychology*, *89*, 672-685.

Gollwitzer, P. M., & Brandstatter, V. (1997). Implementation intentions and effective goal pursuit. *Journal of Personality and Social Psychology*, *73*, 186-199.

Gollwitzer, P. M., Fujita, K., Oettingen, G.(2004). Planning and the implementation of goals. In R. F. Baumeister & K. D. Vohs, (Eds), *Handbook of self-regulation: Research, theory, and applications* (pp. 211-228). New York: Guilford.

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

Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I. An improved scoring algorithm. *Journal of Personality & Social Psychology*, *85*, 197-216.

Hense, R. L., Penner, L. A., & Nelson, D. L. (1995). Implicit memory for age stereotypes. *Social Cognition*, *13*, 399-415.

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

Kawakami, K., Dovidio, J. F., Moll, J., Hermsen, S., & Russin, A. (2000). Just say no (to stereotyping): Effects of training in the negation of stereotypic associations on stereotype activation. *Journal of Personality & Social Psychology*, *78*, 871-888.

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

Liberman, N., & Foerster, J. (2000). Expression after suppression: A motivational explanation of postsuppressional rebound. *Journal of Personality & Social Psychology*, *79*, 190-203.

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

Macrae, C. N., & Bodenhausen, G. V. (2000). Social cognition: Thinking categorically about others. *Annual Review of Psychology*, *51*, 93-120.

Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Jetten, J. (1994). Out of mind but back in sight: Stereotypes on the rebound. *Journal of Personality & Social Psychology*, *67*, 808.

Macrae, C. N., Bodenhausen, G. V., Milne, A. B., Thorn, T. M. J., & Castelli, L. (1997). On the activation of social stereotypes: The moderating role of processing objectives. *Journal of Experimental Social Psychology*, *33*, 471.

Monteith, M. J., Ashburn-Nardo, L., Voils, C. I., & Czopp, A. M. (2002). Putting the brakes on prejudice: On the development and operation of cues for control. *Journal of Personality & Social Psychology*, *83*, 1029-1050.

Nosek, B. A., Greenwald, A. G., Banaji, M. R. (2005). Understanding and using the Implicit Association Test: II. Method variables and construct validity. *Personality & Social Psychology Bulletin*, *31*, 166-180.

Ouellette, J. A., & Wood, W. (1998). Habit and intention in everyday life: The multiple processes by which past behavior predicts future behavior. *Psychological Bulletin*, *124*, 54-74.

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. (2005). Conceptualizing control in social cognition: How executive functioning modulates the expression of automatic stereotyping. *Journal of Personality and Social Psychology*, *84*, 488-503.

Payne, B. K., Jacoby, L. L., & Lambert, A. J. (2004). Memory monitoring and the control of stereotype distortion. *Journal of Experimental Social Psychology*, *40*, 52-64.

Payne, B. K., Jacoby, L. L., & Lambert, A. J. (2005). Attitudes as accessibility bias: Dissociating automatic and controlled processes. In R. R. Hassin, J. S. Uleman, & J. Bargh (Eds.), *The new unconscious* (pp. 393-420). 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.

Plant, E. A., Peruche, B. M., & Butz, D. A. (2005). Eliminating automatic racial bias: Making race non-diagnostic for responses to criminal suspects. *Journal of Experimental Social Psychology*, *41*, 141-156.

Richeson, J. A., Baird, A. A., Gordon, H. L., Heatherton, T. F., Wyland, C. L., Trawalter, S., et al. (2003). An fMRI investigation 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.

Richeson, J. A., & Trawalter, S. (2005). Why do interracial interactions impair executive function? A resource depletion account. *Journal of Personality and Social Psychology*, *88*, 934-947.

Sinclair, L., & Kunda, Z. (1999). Reactions to a Black professional: Motivated inhibition and activation of conflicting stereotypes. *Journal of Personality & Social Psychology*, *77*, 885-904.

Spencer, S. J., Fein, S., Wolfe, C. T., Fong, C., & Dunn, M. A. (1998). Automatic activation of stereotypes: The role of self-image threat. *Personality & Social Psychology Bulletin*, *24*, 1139-1152.

Tipper, S. P. (2001). Does negative priming reflect inhibitory mechanisms? A review and integration of conflicting views. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, *54*, 321-343.

Webb, T. L., & Sheeran, P. (2007). How do implementation intentions promote goal attainment? A test of component processes. *Journal of Experimental Social Psychology*, *34*, 295-302.

Wilson, T. D., & Brekke, N. (1994). Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin*, *116*, 117-142.

Wittenbrink, B., Judd, C. M., & Park, B. (2001). Spontaneous prejudice in context: Variability in automatically activated attitudes. *Journal of Personality & Social Psychology*, *81*, 815-827.

Received February 12, 2008
Revision accepted February 25, 2008