Personality and Social Psychology Bulletin, in press

Person Categorization and Automatic Racial Stereotyping Effects on Weapon Identification

Christopher R. Jones and Russell H. Fazio

The Ohio State University

Word Count: 9249 (main text, references, & notes)

Please address correspondence to:

Christopher R. Jones
Psychology Building
Ohio State University
1835 Neil Avenue
Columbus, OH, 42310
Phone: WORK: 614-688-4611, OTHER: 937-307-2936
Email: jones.2333@buckeyemail.osu.edu

Abstract

Prior stereotyping research provides conflicting evidence regarding the importance of person categorization along a particular dimension for the automatic activation of a stereotype corresponding to that dimension. Experiment 1 replicated a racial stereotyping effect on object identification and examined whether it could be attenuated by encouraging categorization by age. Experiment 2 employed socially complex person stimuli and manipulated whether participants categorized spontaneously or by race. In Experiment 3, the distinctiveness of the racial dimension was manipulated by having Black females appear either in the context of Black males or White females. The results indicated that conditions fostering categorization by race consistently produced automatic racial stereotyping and that conditions fostering non-racial categorization can also eliminate automatic racial stereotyping. Implications for the relationship between automatic stereotype activation and dimension of categorization are discussed.

Key words: stereotyping, categorization, prejudice, automaticity, social cognition.

Person Categorization and Automatic Stereotyping Effects on Object Identification

Stereotypes are cognitive representations of categories of people and their associated attributes. They are an inevitable product of cognitive functioning that allow for prediction of others' actions in the absence of individuating information. Their use conserves cognitive resources (Macrae, Milne, & Bodenhausen, 1994). The extent to which individuals rely on them, however, is a perennial cause for concern. Even when stereotypes do include a "kernel of truth," they are often misleading when applied to a particular individual. Stereotypes also contribute to socially problematic phenomena like prejudice, the outgroup homogeneity effect, and stereotype threat.

As such, much research has sought to examine when stereotype effects can be anticipated and how they might be avoided. This has been particularly true of racial stereotypes, which have been viewed as especially socially problematic. Consider, for example, the finding that individuals (including police officers) can be more likely to mistake[1] harmless objects for weapons when held by or in the presence of Black people relative to White people (Correll, Park, Wittenbrink, & Judd, 2002; Payne, 2001; Plant & Peruche, 2005). Racial stereotypes are also of interest from a social cognitive perspective because race appears to be a particularly salient dimension of person categorization and is afforded a number of cognitive processing advantages. Ito and Urland (2003), for example, measured event-related brain potentials and found that attention was allocated to race at very early stages of processing, even earlier than was observed for attention to the socially and biologically important dimension of gender.

Much of the concern about the influence of stereotypes arises from research on the automaticity of stereotype use, which demonstrates that the effects of stereotypes can

be pernicious. Stereotype activation, the process by which stereotypes are accessed from memory, can be distinguished from stereotype application, when a stereotype influences a thought or behavior towards a member of the stereotyped group. Both of these processes can entail a high degree of automaticity. Consider the finding that stereotypes can be activated by stimuli presented subliminally and can exert downstream consequences of which individuals are unaware, including effects on judgments and behaviors that are not directed at a member of the stereotyped group (Bargh, Chen, & Burrows, 1996; Devine, 1989). It is one thing to knowingly judge another heuristically using stereotypes but quite another for stereotypes to exert an unintended and unnoticed influence.

Although researchers were initially rather pessimistic about the possibility of circumventing automatic stereotype activation in particular (e.g. Bargh, 1999; Devine, 1989), later research revealed that automatic stereotype activation could be reduced in a number of ways. For example, cognitive load can reduce stereotype activation (but can increase stereotype application if load occurs after activation) (Gilbert & Hixon, 1992). Repeatedly negating a stereotype by responding "NO" when it is encountered can also reduce spontaneous stereotype activation (Kawakami, Dovidio, Moll, Hermsin, & Russin, 2000). So can priming a "think different" mindset with creativity words (Sassenberg & Moskowitz, 2007), as can constructing mental images of counterstereotypic exemplars (Blair, Ma, & Lenton, 2001). Perhaps the most encouraging finding was the discovery that simply having a goal to be egalitarian can pre-empt stereotype activation preconsciously, likely because such goals provoke habitual effortful suppression that leads to automatization of an inhibitory response (Moskowitz, Gollwitzer, Wasel, &

Schaal, 1999). Also, specific intentions to think counterstereotypical thoughts when encountering a member of a stereotyped group can successfully reduce unwanted stereotype activation (Stewart & Payne, 2008). The present research examined another factor that might determine when stereotypes become active: the dimension of person categorization. Specifically, we investigated whether the dimension of person categorization would moderate automatic racial stereotyping.

Dependence of Stereotyping on Categorization

Of course, every person can be categorized in many ways, including the salient socio-biological dimensions of gender, race, and age, as well as others. Because stereotypes are attributes associated with categories, activating a particular category is likely to activate an associated stereotype (Dovidio, Evans, & Tyler, 1986) and lead to a processing advantage for stereotype-relevant information (Macrae, Stangor, & Milne, 1994). Zárate and Smith (1990) showed that the latency with which individuals categorized photographs by race, a measure of individual differences in the accessibility of racial categories, predicted the extent to which participants would later make racial stereotype-consistent trait inferences. Macrae, Bodenhausen, and Milne (1995) found that either subliminally priming a particular category label or presenting a cue that made one categorical dimension more salient could moderate stereotype activation. Moreover, in this research, not only did categorizing a person along a particular dimension predict usage of the relevant stereotype, categorizing a person in a particular way produced inhibition of a stereotype corresponding to a competing dimension. Similar findings of moderation of self-stereotyping by self-categorization have also been observed (Shih, Pittinsky, & Ambady, 1999). Thus, it appears that categorization can have both

excitatory effects on stereotypes pertaining to the activated category and inhibitory effects on the accessibility of stereotypes pertaining to competing categories.

Independence of Stereotyping and Categorization

While the preceding discussion suggests a very direct relationship between dimension of categorization and stereotype activation, other studies demonstrate that this relationship is not so straightforward. The key observation is that category and stereotype activation are dissociable. For example, Lepore and Brown (1997) found evidence of automatic racial stereotype application following subliminal *category* primes (i.e. a group label like "Blacks") only for participants relatively high in prejudice (even though knowledge of the cultural stereotype was largely comparable across levels of prejudice), whereas subliminal *stereotype* primes (i.e. specific stereotypical trait words) affected ratings of a target person regardless of level of prejudice.

Even more pertinent to the issue at hand is the possibility that a particular stereotype can be activated even when the corresponding category is not applied. Blair, Judd, and Fallman (2004) identified face stimuli that participants could easily and with high consensus categorize as White or Black. However, these stimuli also varied in the Afrocentricity of features within both racial categories. They found that stimuli with highly Afrocentric features automatically influenced social judgments in a manner associated with the activation of Black stereotypes even when those stimuli were explicitly categorized as White.

In addition to the dissociability of category and stereotype activation, manipulations of dimension of categorization might also prove relatively ineffective in moderating stereotype activation because it might prove difficult to circumvent racial

categorization in the first place, due to its socio-biological salience, early attraction of attention, and potential independence from processing goals (Ito & Urland, 2003). Possibly due to one of these reasons, at least one previous study has found that a simple manipulation of dimension of categorization did *not* prevent racial stereotype activation (Wheeler & Fiske, 2005, Study 2). Although both speculating about the personal preferences of a target (individuation) and processing the target in a non-social manner (i.e. determining whether a dot appeared on the target photo, see also Macrae, Bodenhausen, Milne, Thorn, & Castelli, 1997) eliminated automatic racial stereotype activation, guessing whether that individual was over or under the age of 21 (intended to produce categorization by age) did not.

Therefore, prior research is contradictory regarding the utility of categorization manipulations as a means of controlling stereotyping. To help resolve this discrepancy, the present research examined whether manipulations of dimension of categorization would moderate automatic racial stereotype activation and application.

Automatic Racial Stereotyping and Weapon Bias

Towards this end, we employed the weapon identification procedure, in which faces varying in race are briefly presented before a gun or tool appears. Participants' task is to identify the object as quickly as possible (Payne, 2001). As mentioned, research using this procedure has consistently revealed priming effects of faces varying in race on object identification. In the presence of Black relative to White individuals, participants more easily identify guns, more slowly identify tools, and more frequently tend to respond to tools as if they were guns. Exploration of this effect, sometimes called weapon bias, has clearly supported the contention that it stems from stereotypic racial

associations rather than prejudiced attitudes per se (Correll et al., 2002; Judd, Blair, & Chapleau, 2004; Payne, 2006). Its cognitive underpinning appears to be a commonly held bi-directional association between Blacks and weapons (Eberhardt, Goff, Purdie, & Davies, 2004). Glaser and Knowles (2008) found that as the strength of this association increases, so does the tendency to misidentify objects in a stereotypical fashion, though this relationship was observed only in those not motivated to avoid prejudiced responding.

The weapon identification procedure affords an interesting test of the relationship between dimension of person categorization and automatic stereotyping. Because photos can be used as prime stimuli, membership in multiple categories can be conveyed simultaneously in a naturalistic way. Because exposure to primes is very brief and requires a rapid judgment about an object other than the prime person per se, the weapon bias effect appears highly automatic in the sense of being unintentional and efficient.

This is particularly important because most prior research examining the dimension of categorization and stereotyping used procedures which allowed for relatively more thoughtful stereotyping processes. Often these paradigms use relatively long exposures to multiply categorizable stimuli or create conditions in which participants are motivated to apply a particular stereotype during the judgment task (Macrae et al., 1995; Sinclair & Kunda, 1999). In other research, the dependent variable likely prompts some intentional stereotype application, as when trait inferences about unknown individuals must be made and participants may choose to use stereotypes to complete the task (e.g. Zárate and Smith, 1990). The present research addresses whether more clearly automatic stereotyping is sensitive to manipulation by dimension of

categorization by examining rapid perceptual object identification following racial primes.

In addition to providing information about the generalizability of prior findings relating dimension of categorization and stereotype activation, the role of person categorization may be of interest specifically as a method of eliminating weapon bias. Prior research has demonstrated that practice and expertise in controlling the relevant response are means toward this end (Correll, Park, Judd, Wittenbrink, Sadler, & Keesee, 2007; Plant & Peruche, 2005). Two other manipulations that have reduced weapon bias are forming implementation intentions to think counterstereotypic thoughts (e.g. "Whenever I see a Black person, I will think 'safe'") following Black primes (Stewart & Payne, 2008) and creating experimental conditions that foster the observation that race and the presence/absence of a weapon are uncorrelated, thus rendering race non-diagnostic (Plant, Peruche, and Butz, 2005).

Interestingly, Payne, Lambert, and Jacoby (2002) examined the influence of instructions to either use race or suppress race when identifying the objects. Using race, unsurprisingly, led to weapon bias, but so did attempting to suppress it. Each increased the accessibility of racial stereotypes (in the latter case ironically, see Macrae, Bodenhausen, Milne, & Jetten, 1994). Thus, while categorizing by race but attempting to suppress its automatic effects did not allow participants to avoid weapon bias, perhaps categorization along another dimension would. Experiment 1 examined the relationship between dimension of person categorization and automatic racial stereotyping by examining whether weapon bias would be sensitive to a direct manipulation of dimension

of categorization. Specifically, participants were randomly assigned to categorize primes either by race or by age.

EXPERIMENT 1

METHOD

Participants

Forty-one male (n = 25) and female (n = 16) undergraduates participated for partial course credit. The majority of participants' self-reported ethnicities were Caucasian (n = 35), and none were African-American.

Procedure

Experimental Session. Participants signed up for a half-hour session described as an investigation of their ability to make rapid judgments under time pressure when performing two tasks simultaneously. Upon reaching the laboratory, each participant was seated alone at a computer by a White experimenter. All instructions were provided by the computer. Participants first completed a practice block of weapon identification trials in which they were familiarized with the object categorization task. Following the practice trials, participants were exposed to a manipulation of dimension of categorization. This was followed by three experimental blocks of weapon identification trials.[2] Finally, participants were debriefed and dismissed.

Weapon Identification Procedure. Participants first received instructions for a 32-trial practice block of the weapon identification procedure. They were informed that they would see two photos, a face and an object. The object would be either a gun or a tool. They were told to ignore the face and to simply categorize the object by pressing the

corresponding button on the keyboard labeled “gun” (the “Z” key) or “tool” (the “?” key). “You must respond as quickly as possible!” the instructions continued.

The experimental blocks, like the practice block, involved the following parameters: the face primes appeared for 200 ms, immediately followed by a target object for 100 ms, which was then replaced by a pattern mask. Responses quicker than 200 ms or slower than 450 ms were followed respectively by messages asking participants to wait for the target or to identify it more quickly. Such short response windows are very demanding on participants and typically lead to weapon bias in error patterns rather than response latencies (Payne, 2001). Within each block, order of trials was randomized. Each category of prime (described below) appeared an equal number of times and was followed half the time by a randomly selected gun stimulus and half the time by a randomly selected tool stimulus. The first experimental block of the procedure was 32 trials. The second and third blocks were each 64 trials. Between each block, participants were given the opportunity to rest briefly before continuing the task.

Stimuli. The photos used as primes were 4 Black male and 4 White male faces. These face photos were presented at approximately 5 cm x 4 cm, displayed neutral expressions, and were tightly cropped around the face. All were black-and-white. From each racial category, 2 photos were of young adults previously used by Payne and colleagues (e.g., Stewart & Payne, 2008). The other 2 prime photos from each category were new photos of obviously older adults selected specifically for this study from Internet searches. Prime photos were selected such that participants could easily distinguish 4 relatively “young” faces from 4 relatively “old” faces and could also easily determine race. Thus, primes fell into the categories of young Black, young White, old

Black, and old White. Other stimuli included 8 target photos (4 handguns and 4 hand tools) and a black-and-white static-like rectangular mask, all of which were previously used by Payne and colleagues.

Categorization Manipulation. Following the practice block, instructions informed participants that we were now interested in their ability to perform two tasks at once. They were told that in addition to categorizing the objects, they would be counting how many faces of particular types appeared during the task. Participants were told to click one of four interchangeable boxes on the screen to determine the dimension according to which they would count the faces. Regardless of response, participants were randomly assigned to one of two categorization conditions and received the following message: "You have chosen to count the faces according to their race [age]." The following screen displayed all 8 prime stimuli divided into two rows labeled and arranged either by race or age. Participants were instructed to keep a running mental tally of the number of appearances of both Black versus White [young versus old] faces as they appeared during each block. A paper form was provided on which participants could report their tallies after each block. Such a counting manipulation has previously been demonstrated to be effective in producing categorization along a particular dimension (Olson & Fazio, 2003) and encourages categorization along a particular dimension while leaving information about other dimensions available.

RESULTS AND DISCUSSION

Data Trimming

Following Payne (2001), all trials with responses faster than 100 ms. or slower than 1000 ms. were discarded, 2.1% of the total data points across all participants.

Categorization and Weapon Bias

Preliminary analyses indicated that participant gender did not influence error rates in any simple or interactive fashion with prime race, target object, and categorization condition. Thus, the focal analysis concerned whether weapon bias, signified by a prime race x target interaction on error rates, was moderated by categorization condition. Specifically, a 2 (categorization: age or race, between) x 2 (prime race: Black or White) x 2 (prime age: young or old) x 2 (target object: gun or tool) repeated-measure ANOVA was conducted on proportion of incorrect categorizations. A three-way interaction (Figure 1) between categorization, prime race, and target object emerged: $F(1, 39) = 5.14$, $p < .03$. This interaction can be decomposed by examining prime race x target object interaction in each categorization condition separately. In the race categorization condition, there was a prime x target interaction: $F(1, 21) = 10.26$, $p < .01$. On tool trials, these participants made false "gun" responses on a greater proportion of trials following Black primes ($M = .23$, $SD = .18$) than White primes ($M = .11$, $SD = .11$), $t(21) = 2.9$, $p < .01$. On gun trials, they made false "tool" responses on a *smaller* proportion of trials following Black primes ($M = .12$, $SD = .09$) than White primes ($M = .25$, $SD = .17$), $t(21) = -3.1$, $p < .01$. Thus, racial weapon bias was evident in the categorization by race condition. In the age categorization condition, however, the prime x target interaction was non-significant, $F < 1$. On tool trials, false "gun" responses did not differ as a function of whether the preceding prime was Black ($M = .23$, $SD = .16$) or White ($M = .19$, $SD = .10$), $t(18) = 1.25$, $p = .23$. On gun trials, false "tool" responses also did not differ as a function of whether the prime was Black ($M = .16$, $SD = .08$) or White ($M =$

.17, *SD* = .10), $t(18) < 1$. Thus, no weapon bias was apparent in the age categorization condition.[3]

While categorization by race replicated past findings of a stereotypical weapon bias, encouraging categorization by age sufficed to eliminate this effect, insofar as neither false "gun" responses nor false "tool" differed significantly as a function of the race of the preceding prime.

EXPERIMENT 2

Because Experiment 1 lacked a control condition, it is not apparent whether categorizing by race, categorizing by age, or both were responsible for the observed results. First, we were interested in determining whether racial categorization is so dominant that a ceiling effect would prevent it being increased by drawing attention to race. In a relevant study, Payne, Lambert, and Jacoby (2002) included instructions to "use race" in identifying weapons and tools. This increased weapon bias, but such a manipulation likely directly influences both stereotyping and categorization. Asking participants to use the race of the prime faces to aid in object identification not only directs attention to a dimension of categorization but also essentially requires them to engage in intentional stereotyping. Thus, to determine the effects of mere categorization on automatic stereotyping, Experiment 2 compared a racial categorization condition (again induced by the counting task) to a spontaneous categorization condition.

More importantly, Experiment 2 is also novel with respect to the stimuli employed. Typically, person stimuli used in the weapon identification procedure are selected in such a way that categorization by race is particularly likely. Thus, race is often the only salient dimension along which stimuli vary across trials. Experiment 2

systematically varied race, gender, and occupational status. We were interested in whether participants would spontaneously categorize by race as evidenced by the emergence of weapon bias in the control condition and whether weapon bias would be increased when categorization by race was encouraged.

METHOD

Participants

Seventy-nine Ohio State University undergraduates participated for partial course credit. One participant, an outlier due to the high number of trials eliminated during data trimming, was removed from analyses. This left 78 participants (39 male, 39 female). The majority of participants' self-reported ethnicities were Caucasian (n = 69), and a small proportion were African-American (n = 3)[4].

Procedure

The general procedure followed Experiment 1 closely with the exception of a number of changes to the stimuli and instructions of the weapon identification procedure.

Weapon Identification Procedure. The procedure began again with instructions for a 32 trial practice block during which participants were familiarized with the object identification task. Following instructions constituting the between subjects manipulation of categorization, three blocks of weapon identification trials followed. These blocks were identical in structure to Experiment 1 and involved identical trial parameters. Each of 8 stimuli appeared an equal number of times per block and were followed half the time by a gun and half the time by a tool.

Stimuli. The 8 prime photographs were taken from the work of Fazio and Dunton (1997) and Olson and Fazio (2004). They varied systematically in race (Black vs.

White), gender (male vs. female), and occupational status (high vs. low). Specifically two Black male (minister, gardener), two Black female (businessperson, construction worker), two White male (professor, sanitation worker), and two White female (businessperson, housepainter) photos appeared in a manner otherwise the same as Experiment 1. Weapon and tool stimuli were identical to Experiment 1.

Categorization Manipulation. Following the practice block, participants were informed that the session concerned their ability to complete two tasks at once. Participants were randomly assigned to one of two conditions. Participants in the race categorization condition received instructions consistent with Experiment 1 in which the counting (by race) task was described, the prime photos were presented and distinguished according to race, and a counting form was provided. Participants in the spontaneous categorization condition were told simply to "try to memorize the pictures" that appeared before the gun and tool targets and that a memory test would follow.

RESULTS

Data Trimming

All trials with responses faster than 100 or slower than 1000 ms were again discarded, 1.4% of the total data points across all participants.

Categorization and Weapon Bias

The primary question was again whether weapon bias, a prime by target interaction on error rates, was moderated by categorization condition. Specifically, a 2 (categorization: race or spontaneous, between) x 2 (participant gender[5]: male or female, between) x 2 (prime race: Black or White, within) x 2 (prime gender: male or female,

within) x 2 (prime status: high or low, within) x 2 (target object: gun or tool) mixed ANOVA was conducted on proportion of incorrect categorizations.

The effect of greatest interest, the 3-way categorization condition x prime race x target object interaction, indicates whether racial weapon bias was equally evident in the spontaneous and racial categorization conditions. This interaction was statistically significant, $F(1, 74) = 17.61$, $p < .001$, indicating that weapon bias differed according to categorization condition. This effect was decomposed by conducting separate ANOVA analyzes for the two categorization conditions. In the racial categorization condition, the prime race x target object interaction was significant $F(1, 36) = 17.79$, $p < .001$. Simple effects show that this interaction took the form of racial weapon bias. On tool trials, these participants made more false "gun" errors following Black primes ($M = .26$, $SD = .15$) than White primes ($M = .15$, $SD = .11$), $t(37) = 3.50$, $p = .001$. On gun trials, these participants made fewer false "tool" errors following Black primes ($M = .13$, $SD = .10$) than White primes ($M = .24$, $SD = .15$), $t(37) = 5.15$, $p < .001$. In the spontaneous categorization condition, however, the prime race x target object interaction was non-significant, $F(1, 38) < 1$. Neither errors on gun trials nor on tool trials differed as a function of prime race, p's $> .24$.

This 3-way interaction indicating racial weapon bias following racial categorization but not spontaneous categorization was qualified by a 4-way interaction with prime status $F(1, 74) = 4.02$, $p < .05$ (Figure 2). However, this higher-order interaction does not require modification of the previous interpretation. Importantly, the categorization condition x prime race x target object interaction obtained for both high status primes, $F(1, 74) = 19.42$, $p < .001$, and the low status primes, $F(1, 74) = 9.68$, $p <$

.01. In both cases a pattern of errors consistent with racial weapon bias emerged in the racial categorization condition but not in the spontaneous categorization condition. Thus, while racial weapon bias strongly determined patterns of error in the racial categorization condition, it was not evident under spontaneous categorization.

Interestingly, a categorization condition x prime gender x target object interaction was also observed: $F(1, 74) = 16.18$, $p < .001$ (Figure 3). In the racial categorization condition, the prime gender x target object interaction was non-significant, $F(1, 36) < 1$. In the spontaneous categorization condition, the prime gender x target object interaction was significant, $F(1, 38) = 36.94$, $p < .001$. On tool trials, these participants made more false "gun" errors after male primes ($M = .17$, $SD = .10$) than after female primes ($M = .11$, $SD = .09$), $t(39) = 4.90$, $p < .001$. On gun trials, these participants made fewer false "tool" errors after male primes ($M = .11$, $SD = .09$) than after female primes ($M = .15$, $SD = .11$), $t(39) = 3.48$, $p = .001$. This 3-way interaction was not qualified by higher-order interactions.

DISCUSSION

In summary, drawing attention to race appeared to increase racial stereotype activation and racial weapon bias beyond that observed in a spontaneous categorization condition. In fact, no racial weapon bias was evident given this set of prime stimuli unless attention was drawn to race. Given sufficiently complex stimuli, the racial dimension may not always dominate categorization. Indeed, when categorization was spontaneous, object identification was instead influenced by gender of the prime. Photos of males cued "guns" more strongly than did photos of females. It is interesting that racial weapon bias did not occur in this condition. Presumably, this is because primes

were often categorized by gender and perhaps also to a significant extent by occupation, which, readers should note, is not necessarily the same as categorizing by occupational status, nor would it necessarily be apparent in effects on object identification. Typically, research on person categorization has focused on the "big three" of age, race, and gender. Much less attention has been paid to more fundamentally social dimensions of categorization like occupation. Future research should address these dimensions more thoroughly because it is likely that people use them regularly in their daily lives. Often, for example, occupation might be expected to drive categorization because occupations are associated with social norms, roles, and scripts that regularly dictate behavior toward the target person in a very specific and explicit way that is absent for a dimension like race. Regardless, Experiment 2's primary finding was that encouraging categorization along the dimension of race increased racial weapon bias. This produced a racial stereotyping effect on object identification for both that was not evident in the spontaneous categorization condition.

EXPERIMENTS 3a AND 3b

Experiments 1 and 2 involved a strong manipulation of categorization in which at least one condition absolutely required categorization along a particular dimension. In Experiment 3, we were interested in whether a more subtle manipulation of categorization could moderate weapon bias. To this end, we constructed a version of the weapon identification procedure in which critical trials involved Black female primes. These primes appeared in one of two contexts: the majority of trials involved either Black male primes or White female primes. Previous research indicates that one determinant of dimension of person categorization is distinctiveness (McGuire, McGuire, Child, &

Fujioka, 1978; Taylor, Fiske, Etcoff, & Ruderman, 1978). For example, in a classroom full of White students, the race of a Black student would be particularly salient to everyone. In this case, when a Black female appears in the context of mostly Black males, her gender should be relatively salient (but not her race), whereas when a Black female appears in the context of mostly White females, her race (but not her gender) should be relatively salient. Experiments 3a and 3b were nearly identical experiments conducted in consecutive academic quarters with minor procedural variations and are thus presented together. Two samples were collected in order to replicate an unexpected effect of participant gender that appeared.

METHOD

Participants

3a: Eighty undergraduates participated for partial course credit. One participant, an outlier due to the high number of trials eliminated during data trimming, was excluded from analyses. This left 79 participants (40 male, 39 female). The majority of participants' self-reported ethnicities were Caucasian (n = 62), and a small proportion were African-American (n = 6). 3b: One hundred twenty-five undergraduates participated for partial course credit (54 male, 72 female; 96 Caucasian, 6 African-American).

Procedure

The general procedures followed Experiments 1 and 2 closely with the exception of a number of changes to the stimuli and instructions of the weapon identification procedure. Participants were randomly assigned to one of two conditions determined by

the type of prime faces constituting the distinctiveness manipulation in the weapon identification procedure.

Weapon Identification Procedure: Experiment 3a. Again, participants were first familiarized with the object identification task via a practice block. Next, all participants received instructions which asked them to try to memorize the faces for a memory test that would supposedly follow (as in the spontaneous categorization condition of Experiment 2).

Participants then completed three blocks of the weapon identification procedure. Each block involved 68 trials. These blocks were arranged in a fixed-random order of 14 sub-blocks. In each sub-block either 2, 4, or 6 context trials (involving Black male or White female primes depending on condition) preceded one critical trial involving a Black female prime. We varied the number of context trials that preceded the critical Black female prime trials in order to reduce the predictability with which a Black female was presented. Five sub-blocks involved 2 context trials, 5 sub-blocks involved 4 context trials, and 4 sub-blocks involved 6 context trials. This adds to 54 context trials and 14 critical trials per block. Black female prime trials were therefore distinctive, as they were always preceded by at least two context trials and in sum composed only about 21% of trials whereas the other 79% of trials were one of two types of context trials. All trials involved a randomly selected target object, either a gun or tool. The timing parameters of the trials were the same as Experiment 1 and 2.

Weapon Identification Procedure: Experiment 3b. Experiment 3b was very similar to 3a. The only differences follow. Each block still involved 68 trials arranged in a fixed-random order of 14 sub-blocks of context trials (Black male or White female

primes depending on condition) preceding one critical trial involving a Black female prime. Each sub-block involved either 3, 4, or 5 context trials preceding a critical trial. This change was made because examination of Experiment 3a data suggested that no effect of context was evident on critical trials preceded by only 2 context trials. Six sub-blocks involved 3 context trials, 4 sub-blocks involved 4 context trials, and 4 sub-blocks involved 5 context trials. For all categories of primes, a gun appeared half of the time and a tool appeared half of the time.

Stimuli. The prime stimuli involved 8 photos of Black females, 8 photos of Black males, and 8 photos of White females. All photos were black-and-white, closely cropped around the face, and neutral in facial expression. Their presentation was otherwise the same as in the previous studies.

RESULTS

Data Trimming

All trials with responses faster than 100 or slower than 1000 ms were discarded, 1.4% of the total data points across all participants in Experiment 3a and 0.5% in Experiment 3b.

Categorization and Weapon Bias: Black Female Trials

A 2 (context: Black male or White female, between) x 2 (participant gender: male or female, between) x 2 (target object: gun or tool, within) x 2 data sample (Experiment 3a or Experiment 3b, between) mixed ANOVA was conducted on proportion of incorrect categorizations on Black female trials. If distinctiveness determined categorization and thus stereotyping, we would anticipate a context condition x target object interaction. Specifically, a Black male context should facilitate categorization by gender, while a

White female context should facilitate categorization by race (evident in error rates consistent with racial weapon bias). An unexpected interaction between participant gender, context condition, and target object was obtained, $F(1, 196) = 4.39$, $p < .04$. Importantly, this interaction was not qualified by a higher-order interaction with data sample ($F < 1$), nor did this factor produce any main or interactive effects in this analysis indicating the pattern of data replicated across the two experiments. For men, context condition interacted with target object, $F(1, 90) = 7.64$, $p < .01$, while for women it did not $F(1, 106) = .26$, $p > .60$ (see Figure 4). For male participants, on tool trials, false "gun" errors were more frequent when context encouraged categorization by race ($M = .18$, $SD = .12$) than when context encouraged categorization by gender ($M = .13$, $SD = .12$), $t(92) = 2.14$, $p < .05$. False "tool" errors did not differ by categorization context for male participants, $t < 1$. Notably, when context encouraged categorization by gender, male participants' error rates on gun and tool trials were equivalent, M's = .13, $t < 1$.

However, for female participants, false "gun" error rates were equivalent whether context encouraged categorization by race ($M = .21$, $SD = .16$) or by gender ($M = .18$, $SD = .13$), $t(108) = 1.15$, $p = .25$. False "tool" error rates were also equivalent whether context encouraged categorization by race ($M = .13$, $SD = .11$) or by gender ($M = .11$, $SD = .08$), $t(108) = 1.15$, $p = .25$. Thus, context did not influence female participants; they mistook tools for guns more often than guns for tools following Black female primes in both context conditions, p's $< .01$.

Categorization and Weapon Bias: Context Trials

A second mixed-design ANOVA with the same factors described in the previous section was conducted on the proportion of errors on Black male and White female

context trials. Here, a context condition x target object interaction did appear, $F(1, 196) = 11.17$, $p = .001$. Though false "gun" errors exceeded false "tool" errors for both Black male and White female targets (p's < .01), this interaction indicates that, consistent with what would be anticipated on the basis of both racial and gender stereotyping, this difference was larger following Black male primes ($M_{difference} = .07$, $SD = .11$) than White female primes ($M_{difference} = .03$, $SD = .07$). In this case, participant gender did not moderate the context condition x target object interaction, $F(1, 196) = 2.44$, $p = .12$, although the data pattern trended in the direction of females showing a larger effect than males. Again, there were no effects of data sample, indicating that these results replicated across Experiments 3a and 3b.

DISCUSSION

Experiments 3a and 3b showed that a much more subtle contextual manipulation of categorization can also moderate weapon bias. However, this effect emerged only for male participants. This gender effect might have emerged for various reasons. One way of explaining such an effect would posit that either the accessibility of the relevant racial stereotype or the propensity to categorize by race is stronger for females than males. This possibility does not appear very likely, however. Neither the results of the other experiments reported here nor the literature more generally favors this interpretation. Another possibility is that the manipulation functioned similarly for both participant genders, but that feature-based, categorization-independent racial stereotyping occurred to a greater extent for female than male participants. We are not aware, however, of any particular reason this would be the case.

In our view, the most plausible explanation is that the context manipulation, particularly the condition designed to foster categorization by gender, produced an effect only for male participants because the relevant gender dynamics led it to be more powerful for males. In the condition fostering categorization by race, error rates on Black female trials were highly similar for male and female participants. Both displayed a higher proportion of false "gun" errors than false "tool" errors consistent with racial categorization. Therefore, it appears that the distinctiveness of race on these Black female trials did not differ between male and female participants. However, male and female participants differed in how they responded to Black female primes within a context involving images of Black males. Error rates suggested that females continued to categorize by race, whereas male participants categorized the Black female primes more by gender. Thus, it appears that, relative to females, males were more easily led to focus on the gender of the distinctive opposite-sex female primes, finding that dimension relatively more interesting and attention grabbing. Recall that the Experiment 3 prime stimuli were facial close ups. It may be the case that the dimension of physical attractiveness (itself linked closely to gender) was particularly salient for the male participants when a Black female prime was presented. Such salience may have dampened tendencies to categorize by race more so for male than for female participants. Distinctiveness is known to play an important role in the development of illusory correlations (e.g Hamilton & Gifford, 1967; Strossner & Plaks, 2001). Our speculations suggest that distinctiveness may play a similar role in categorization processes. However, the extent and power of any such contextually-based distinctiveness may itself be greater when it aligns with pre-existing tendencies.

In any case, Experiments 3a and 3b demonstrate that a relatively subtle manipulation of dimension of person categorization can eliminate patterns of error in object identification consistent with racial bias, although this effect was only observed for male participants. Nonetheless, this exception illustrates the fundamental importance of the person categorization process we have highlighted.

GENERAL DISCUSSION

Three experiments examined whether manipulations of the dimension of person categorization would moderate automatic racial stereotyping. Each experiment employed the weapon identification procedure (Payne, 2001), in which participants identify objects preceded by prime faces as guns or tools. Each experiment included a manipulation intended to affect the dimension of person categorization to determine whether weapon bias would occur, an effect in which racial stereotypes cause guns to be more easily identified than tools following Black primes compared to White primes. Experiment 1 replicated this effect when a manipulation strongly encouraged categorization by race. This effect was eliminated by a manipulation that strongly encouraged categorization by age. Experiment 2 employed complex person stimuli and demonstrated that a strong manipulation of attention to race increased racial stereotype activation following both Black and White primes. Experiments 3a and 3b found that a relatively more subtle manipulation of dimension of person categorization involving distinctiveness could influence automatic racial stereotyping of Black female primes, at least for some participants.

The clearest conclusion provided by these studies is that conditions fostering attention to race led to automatic racial stereotype activation. This was evident in all

three experiments. The research also suggests that automatic racial stereotyping can be attenuated by directing attention to a dimension other than race. In Experiment 1, the tendency for prime race to influence the frequency with which individuals mistake relatively harmless tools for dangerous weapons was not observed when attention was directed to the age of the primes. In Experiment 3, this weapon bias did not occur for Black female primes when attention was directed to gender, albeit only for male participants. Moreover, in Experiment 2, complex multiply-categorizable stimuli that simply allowed for attention to be spontaneously directed to dimensions other than race yielded no signs of a racial bias.

Experiment 2 is also noteworthy because it provides clear evidence that that weapon identification was driven not only by a Black stereotype but also a White stereotype and gender stereotypes. First, categorization by race led to greater false "gun" responses and fewer false "tool" responses following Black primes compared to the spontaneous categorization condition. Less anticipated was that this strong manipulation which forced categorization of White primes by race produced the opposite pattern of responses, more false "tool" errors and fewer false "gun" errors, a pattern that also did not occur in the spontaneous categorization condition. Second, guns appeared to be associated with males more than with females. These findings indicate the ease with which weapons are identified is influenced by multiple stereotypes. The effect of racial categorization of White primes suggests that in research on weapon bias the White prime conditions do not serve as mere control conditions against which to compare the effects of Black stereotypes. Instead, responses to White primes can themselves be driven by stereotypes. However, the exact nature of this effect is unclear, and it is not apparent that

this is a function of a well-formed stereotype about White people that our mostly White participants hold and activate spontaneously. Rather, it may only arise in a comparative, interracial context. Regardless, these studies more clearly than others suggest that weapon bias is also driven by responses to Whites and therefore might even be thought of as having a component that represents a sort of White privilege.

The gender effect may seem puzzling insofar as while it may not be surprising that guns were more associated with men than women, one might also expect tools to be. However, in this procedure and likely in general, guns are much more salient objects than tools. In fact, it is not uncommon in the weapon identification procedure and related methodologies to observe main effects of target object such that participants are faster to identify guns than tools and make fewer errors on gun trials. Further, the specific gun photos are all relatively similar photos that are clearly handguns, while the particular tools differ more substantially and are in some cases comparatively difficult to specifically identify. For these reasons to some extent many participants may effectively be performing a "gun" versus "not gun" judgment task, which is not problematic because the use of tools is arbitrary in the first place. If the assertion that task performance is driven more by responses to guns than tools is accepted, Experiment 2's prime gender effects within the spontaneous categorization condition are unsurprising.

These results of these experiments are informative regarding prior inconsistencies in the literature of categorization and automatic racial stereotyping. First, they argue against strong interpretations of earlier research findings in which manipulations of dimension of categorization did not attenuate racial stereotyping. We found that both a strong manipulation drawing attention away from race and a relatively subtle one could

do so. Further research might explore the boundary conditions of such effects. The ease of categorizing the perceived individual along the non-racial dimension of categorization is an obvious candidate. To the extent that the non-racial dimension is itself salient, automatic racial stereotype activation will be less likely. Second, these results reaffirm the importance of the category-stereotype link. Drawing attention to race very consistently produced substantial weapon bias. These results do not, however, contradict earlier findings (e.g. Blair, Judd, & Fallman, 2004) that stereotype activation need not be mediated by the activation of the corresponding category. There may be multiple pathways to the activation of any given stereotype, but fostering or inhibiting potentially relevant categorizations clearly affects the likelihood of such stereotype activation.

Ultimately, automatic racial stereotyping likely can be explained largely by well-established principles of knowledge activation. A stereotype might be activated by activation spreading through any concept with which it is associated in memory. Though associations between a given category and stereotype may vary in strength across individuals, it stands to reason that activating a category should increase the likelihood of activation of a stereotype that corresponds to it, and the data from this research were strongly consistent with this idea. Other associates of stereotypes, such as prototypical features and traits, essentially the content of a stereotype itself, might also be associated with a particular stereotyped attribute and could activate it in a manner independent of category activation, and the data from this research do not contradict this idea. On the other hand, encouraging categorization by dimensions other than race appeared to reduce racial stereotyping, indicating that interventions against racial (and likely other) stereotype effects might fruitfully focus on categorization. One question this raises

concerns whether the effects we observed would generalize from the type of horizontal manipulation of dimension of person categorization used in the present research to a vertical manipulation of categorization such as categorizing targets in a more inclusive way as, for example, Americans, or even simply as people. Such an approach has received considerable attention in the domain of prejudice (Brewer, 2007; Gaertner & Dovidio, 2000). It may be similarly useful to examine its implications for the purely cognitive element of automatic stereotype use per se. This is a potential avenue for further inquiry.

REFERENCE LIST

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

Bargh, J. A., Chen, M., & Burrows, L. (1996). Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action. *Journal of Personality and Social Psychology*, 71, 230-244.

Blair, I. V., Judd, C. M., & Fallman, J. L. (2004). The automaticity of race and Afrocentric facial features in social judgments. *Journal of Personality and Social Psychology*, *87*, 763-778.

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

Brewer, M. B. (2007). The social psychology of intergroup relations: Social categorization, ingroup bias, and outgroup prejudice. In A. W. Kruglanski & E. T. Higgins (Eds.), *Social psychology: Handbook of basic principles*, 2nd ed (pp. 695-715*)*. New York: Guilford.

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

Correll, J., Park, B., Judd, C. M., Wittenbrink, B., Sadler, M. S., & Keesee, T. (2007). Across the thin blue line: Police officers and racial bias in the decision to shoot. *Journal of Personality and Social Psychology*, 92, 1006-1023.

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

Dovidio, J. F., Evans, N., & Tyler, R. B. (1986). Racial stereotypes: The content of their cognitive representations. *Journal of Experimental Social Psychology*, *22*, 22-37.

Eberthardt, J. L., Goff, P. A., Purdie, V. J., & Davies, P. G. (2004). Seeing Black: Race, crime, and visual processing. *Journal of Personality and Social Psychology*, 87, 876-893.

Fazio, R. H. & Dunton, B. C. (1997). Categorization by race: The automatic and controlled components of racial prejudice. *Journal of Experimental Social Psychology*, *33*, 451-470.

Gaertner, S. L. & Dovidio, J. F. (2000). *Reducing intergroup bias: The common ingroup identity model.* Philadelphia: Psychology Press.

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.

Glaser, J. & Knowles, E. D. (2008). Implicit motivation to control prejudice. *Journal of Experimental Social Psychology*, *44*, 164-172.

Hamilton, D. L., & Gifford, R. K. (1976). Illusory correlation in intergroup perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology*, *12*, 392-407.

Henry, P. J. & Sears, D. O. (2002). The Symbolic Racism 2000 Scale. *Political Psychology, 23*, 253-283.

Ito, T. A. & Urland. G. R. (2003). Race and gender on the brain: Electrocortical measures of attention to the race and gender of multiply categorizable individuals.

Journal of Personality and Social Psychology, 85, 616-626.

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

Lepore, L. & Brown, R. (1997). Category and stereotype activation: Is prejudice inevitable? *Journal of Personality and Social Psychology*, 72, 275-287.

Macrae, C. N., Bodenhausen, G. V., & Milne, A. B. (1995). The dissection of selection in person perception: Inhibitory processes in social stereotyping. *Journal of Personality and Social Psychology*, *69*, 397-407.

Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Jetten, J. (1994). Out of mind but back in sight: Stereotypes on the rebound. *Journal of Personality and Social Psychology*, *67*, 808-817.

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, 471-489.

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.

Macrae, C. N., Stangor, C., & Milne, A. B. (1994). Activating social stereotypes: A functional analysis. *Journal of Experimental Social Psychology*, *30*, 370-389.

McGuire, W. J., McGuire, C. V., Child, P., & Fujioka, T. (1978). Salience of ethnicity in the spontaneous self-concept as a function of one's ethnic distinctiveness in the social environment. *Journal of Personality and Social Psychology, 36*, 511-520.

Moskowitz, G. B., Gollwitzer, P. M., Wasel, W., & Schaal, B. (1999). Preconscious control of stereotype activation through chronic egalitarian goals. *Journal of Personality and Social Psychology*, 77, 167-184.

Olson, M. A., & Fazio, R. H. (2003). Relations between implicit measures of prejudice: What are we measuring? *Psychological Science, 14,* 636-639.

Olson, M. A., & Fazio, R. H. (2004). Trait inferences as a function of automatically activated racial attitudes and motivation to control prejudiced reactions. *Basic and Applied Social Psychology*, *26*, 1-11.

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. (2006). Weapon bias: Split-second decisions and unintended stereotyping. *Current Directions in Psychological Science,* 15, 287-291.

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.

Payne, B. K., Shimizu, Y., & Jacoby, L. L. (2004). Mental control and visual illusions: Toward explaining race-based weapon misidentifications. *Journal of Experimental Social Psychology*, *41,* 36-47.

Plant, E. A. & Devine, P. G. (1998). Internal and external motivation to respond without prejudice. *Journal of Personality & Social Psychology, 75,* 811-832.

Plant, E. A. & Peruche, B. M. (2005). The consequences of race for police officers' responses to criminal suspects. *Psychological Science*, *16*, 180-187.

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

Sassenberg, K. & Moskowitz, G. B. (2005). Don't stereotype, think different! Overcoming automatic stereotype activation by mindset priming. *Journal of Experimental Social Psychology,* 41, 506-514.

Shih, M., Pittinsky, T. L., & Ambady, N. (1999). Stereotype susceptibility: Identity salience and shifts in quantitative performance. *Psychological Science, 10,* 80-83.

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.

Stewart, B. D. & Payne, B. K. (2008). Bringing automatic stereotyping under control: Implementation intentions as efficient means of thought control. *Personality and Social Psychology Bulletin, 34,* 1332-1345.

Stroessner, S. J. & Plaks, J. E. (2001). Illusory correlation and stereotype formation: Tracing the arc of research over a quarter century. In G. B. Moskowitz (Ed.), *Cognitive social psychology: The Princeton Symposium on the Legacy and Future of Social Cognition* (pp. 247-259). Mahwah, NJ: Erlbaum.

Wheeler, M. E., & Fiske, S. T. (2005). Controlling racial prejudice: Social-cognitive goals affect amygdala and stereotype activation. *Psychological Science,* 56-63.

Taylor, S. E., Fiske, S. T., Etcoff, N. L., & Ruderman, A. J. (1978). Categorical and contextual bases of person memory and stereotyping. *Journal of Personality and Social Psychology, 36,* 778-793.

Zárate, M. A. & Smith, E. E. (1990). Person categorization and stereotyping. *Social Cognition*, 8, 161-185.

NOTES

1. Although the term “mistake” and its variants will repeatedly be used to describe this sort of error in object identification, its meaning should be clarified. Perceivers are mistaken insofar as an incorrect response is given. This should not be taken to mean that they literally misperceive the object, at least in anything but an extremely transient fashion. The available data are more consistent with the possibility that perceivers generally correctly identify the object but fail to appropriately execute the correct response due to stereotype-induced response bias (Payne, Shimizu, & Jacoby, 2004).

2. After the weapon identification procedure, participants also completed two individual difference scale measures after the weapon identification procedure, the symbolic racism scale (Henry & Sears, 2002), and a measure of motivation to control prejudice against Black people (Plant & Devine, 1998). However, because the few effects that did obtain with these scale measures were weak and did not replicate across studies, results concerning these individual difference variables are not reported.

3. One might expect that while all prime faces might be easily categorized by race, the older prime faces might especially lend themselves to categorization by age, thus producing a 4-way interaction in which the effect of categorization on weapon bias was more evident for the older Black and White faces. Although in this analysis the 4-way interaction was not statistically significant, $F(1, 39) = 1.30$, $p = .26$, it should be noted that N is quite small here for such a test, and inspection of the data did reveal a trend such that the categorization manipulation did appear somewhat more influential for the old primes than the young ones. Therefore, the data suffice to demonstrate that directing attention to a dimension other than race can moderate weapon bias, but the possibility

that this is most or even only true for individuals who are particularly easily categorized along a dimension other than race should not be dismissed.

4. The exclusion of African-American participants from analyses would not affect the results, i.e. the general form and statistical significance of the effects of categorization condition on weapon bias. The same holds true for Experiment 3.

5. Participant gender is retained as a design factor in Experiment 2 because, unlike Experiment 1, it did produce statistically significant effects in the ANOVA analysis, though these did not involve the focal analyses concerning the effect of categorization condition on weapon bias and are thus not discussed.

AUTHORS' NOTE

Christopher R. Jones, Department of Psychology, The Ohio State University; Russell H. Fazio, Department of Psychology, The Ohio State University. The authors thank Brandon Stewart, Chris Loersch, and the Social Cognition Research Group for their contributions. Please address correspondence to: Christopher Jones, Psychology Building, Ohio State University, 1835 Neil Avenue, Columbus, OH 43210-1287. E-mail: jones.2333@buckeyemail.osu.edu

FIGURE CAPTIONS

Figure 1: Interaction between categorization condition, prime race, and target object predicts error rate. Error bars indicate + and - one standard error of the mean.

Figure 2: Interaction between categorization condition, prime race, prime status, and target object predicts error rate. Abbreviations: Spont = spontaneous, and Cat = categorization condition. Error bars indicate + and - one standard error of the mean.

Figure 3: Interaction between categorization condition, prime gender, and target object condition predicts error rate. Error bars indicate + and - one standard error of the mean.

Figure 4: Combined data from Experiments 3a and 3b. Interaction between participant gender, context condition, and target object predicts error rate on critical trials involving Black female primes. Error bars indicate + and - one standard error of the mean.

Figure 1.

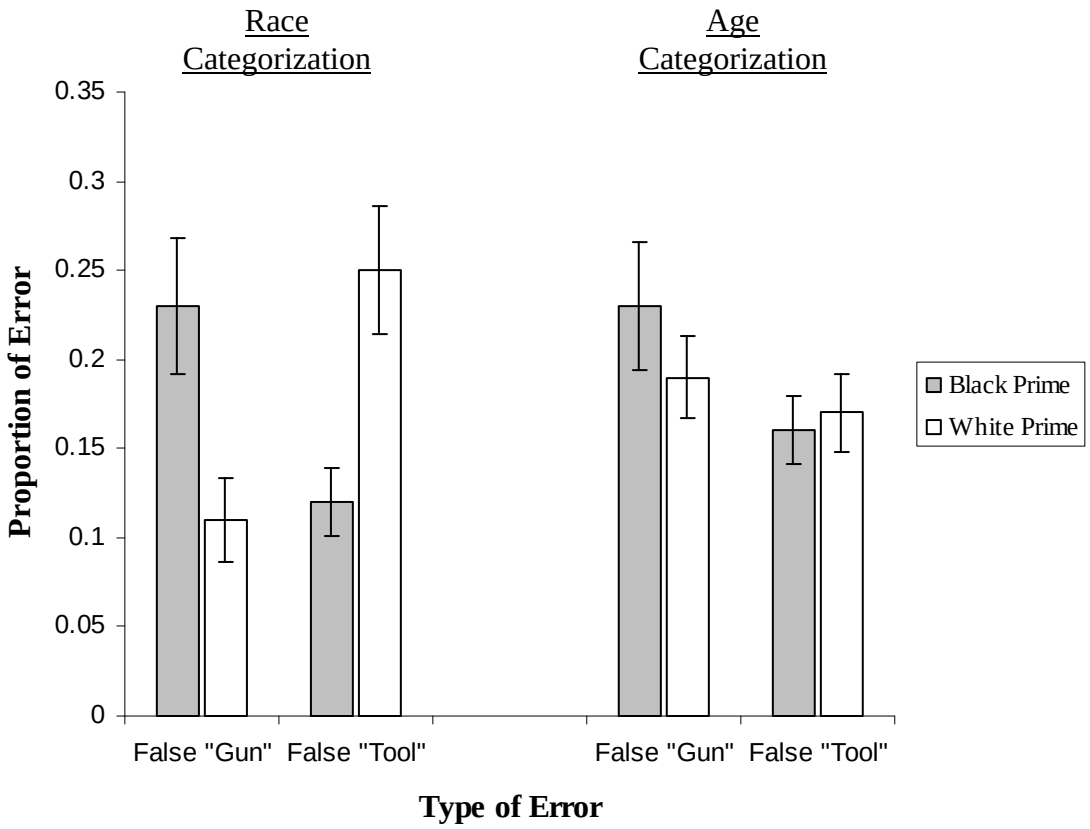
Race Categorization
Age Categorization
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
Proportion of Error
Black Prime
White Prime
False "Gun"
False "Tool"
False "Gun"
False "Tool"
Type of Error

Figure 2.

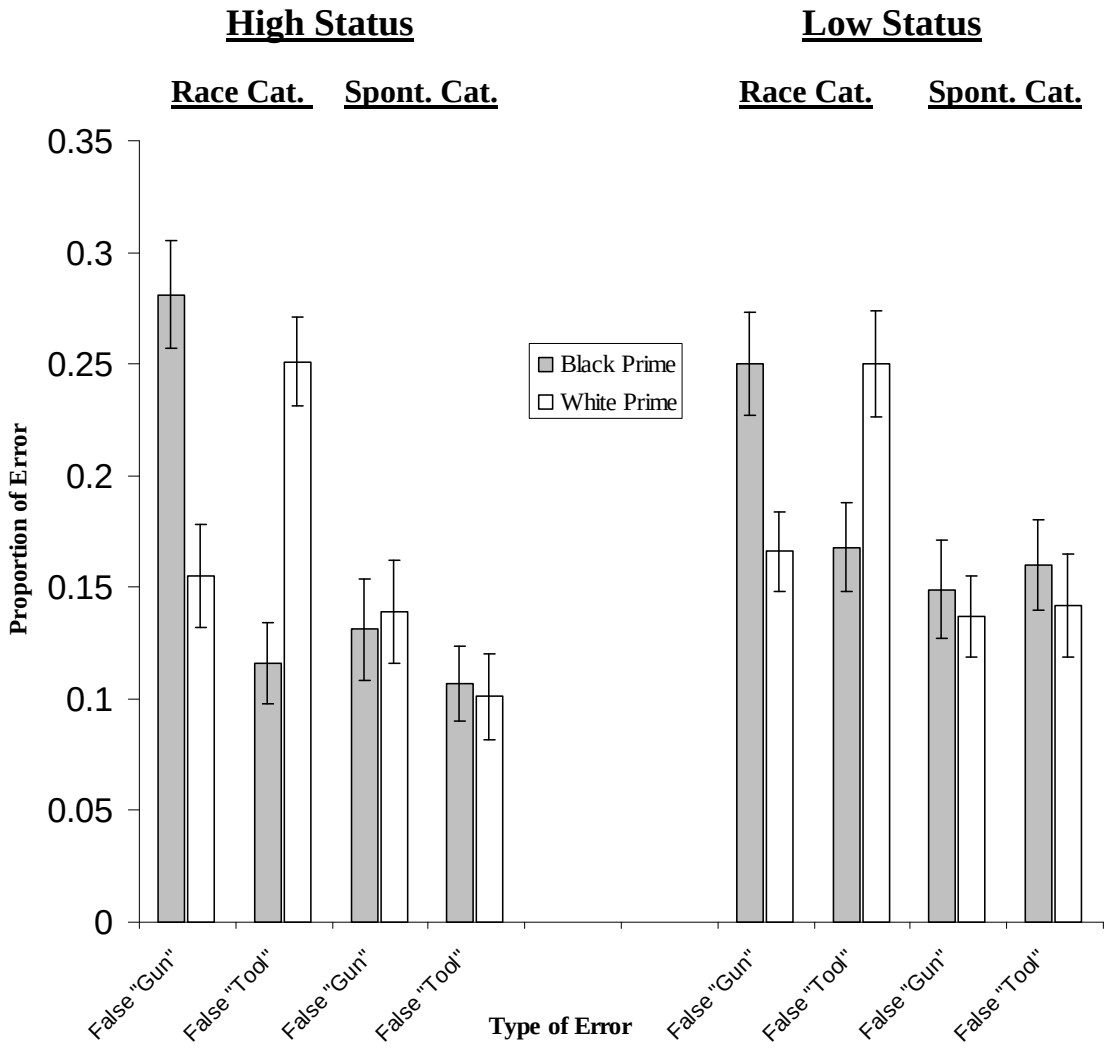
High Status
Low Status
Race Cat.
Spont. Cat.
Race Cat.
Spont. Cat.
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
Proportion of Error
Black Prime
White Prime
False "Gun"
False "Tool"
False "Gun"
False "Tool"
Type of Error
False "Gun"
False "Tool"
False "Gun"
False "Tool"

Figure 3.

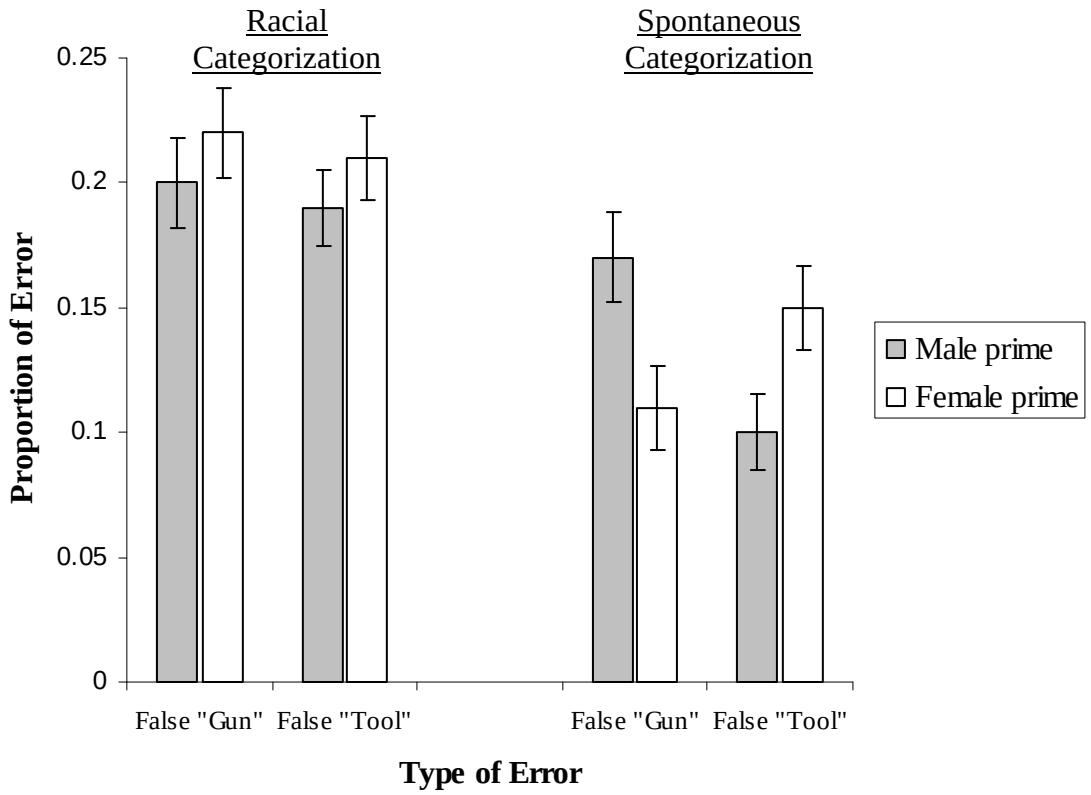
Racial Categorization
Spontaneous Categorization
0.25
0.2
0.15
0.1
0.05
0
Proportion of Error
False "Gun"
False "Tool"
False "Gun"
False "Tool"
Type of Error
Male prime
Female prime

Figure 4.

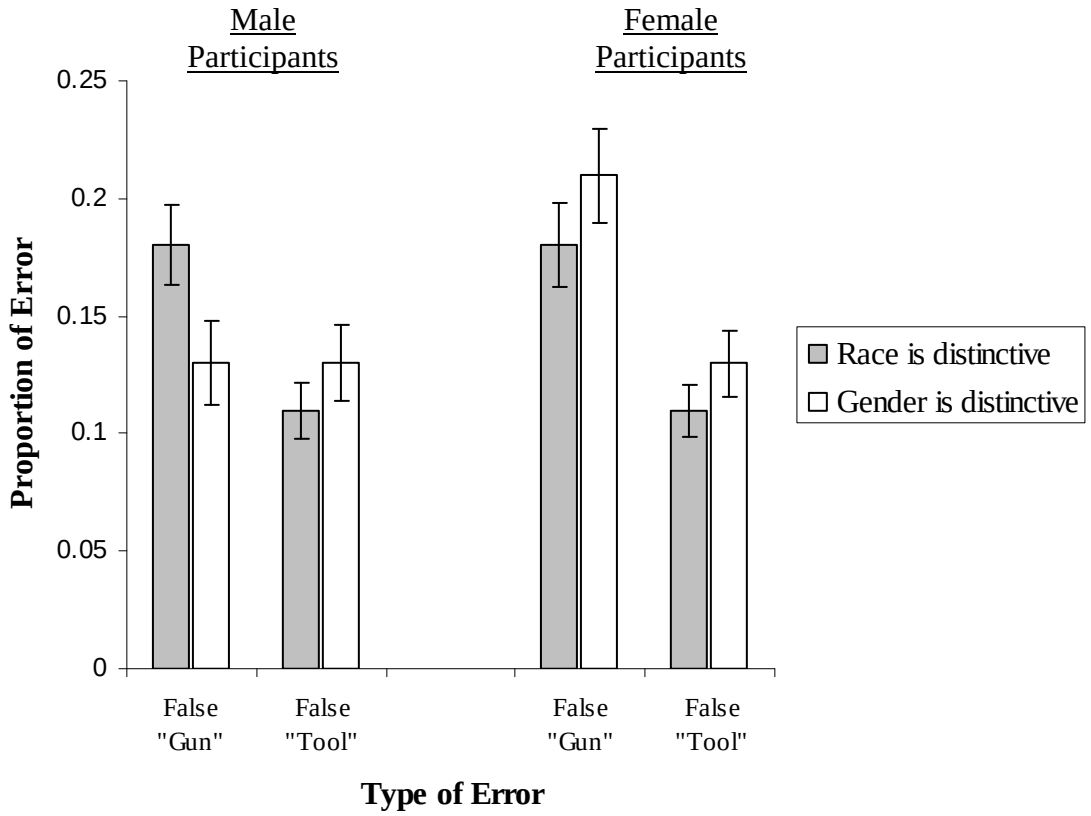
Male Participants
Female Participants
Proportion of Error
0.25
0.2
0.15
0.1
0.05
0
False "Gun"
False "Tool"
False "Gun"
False "Tool"
Type of Error
Race is distinctive
Gender is distinctive