



Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

journal homepage: www.elsevier.com/locate/jesp

Reports

Danger stereotypes predict racially biased attentional allocation[☆]Nicole C. Donders^{*}, Joshua Correll, Bernd Wittenbrink

University of Chicago, Department of Psychology, 5848 South University Avenue, G-117 Chicago, IL 60637, USA

ARTICLE INFO

Article history:

Received 7 December 2007

Revised 9 April 2008

Available online 22 April 2008

Keywords:

Race

Stereotyping

Prejudice

Fear-conditioning

Attentional bias

Dot-probe

ABSTRACT

The current research investigates whether threat-relevant associations have specific implications for attentional allocation, over and above the effects of other category-based associations. Using a modified dot-probe task [Koster, Crombez, Verschuere, & DeHouwer (2004)], we separately measured attentional capture and attentional holding by Black compared to White faces. Black-danger associations significantly predicted the extent to which Black faces captured attention faster than White faces. Black-danger stereotypes also marginally predicted the extent to which Black faces held attention longer than White faces. These effects remained significant when controlling for the effects of other (danger-irrelevant) stereotypes and prejudice, and neither danger-irrelevant stereotypes nor prejudice predicted racially biased attentional allocation. We posit that societal stereotypes linking Blacks with danger lead Black faces to function as fear-conditioned stimuli, biasing attention.

© 2008 Elsevier Inc. All rights reserved.

Introduction

Stereotypes often form a complex, multifaceted set of associations that includes both positive and negative components. For example, traits like poor, lazy, and athletic are often associated with Black people in explicit and implicit measures (Devine, 1989; Dovidio & Gaertner, 1986; Wittenbrink, Judd, & Park, 1997). In addition, research has focused heavily on general negativity or *prejudice* toward Blacks, which is typically treated as independent of specific semantic content. Highly prejudiced individuals associate Blacks with any negative concept, even if that concept is not particularly relevant to the Black stereotype. For example, exposure to Black faces can facilitate responses to negative but nonstereotypic words like “poison” or “cancer” (Fazio, Jackson, Dunton, & Williams, 1995; Greenwald, McGhee, & Schwartz, 1998).

A prevalent component of the Black stereotype relates specifically to *danger* or *threat*. Many studies have demonstrated that stereotypes associating Blacks with danger can dramatically affect judgment and behavior. For example, perceivers interpret an actor’s ambiguous behavior as more hostile when that actor is Black than when he is White (Duncan, 1976; Sagar & Schofield, 1980). Devine (1989) subliminally primed some participants with traits that were stereotypic of Blacks but unrelated to the concept of violence. These participants later judged an ambiguous social target as

more hostile than did subjects who were not primed with Black stereotypes, suggesting that activating the social category facilitated perceptions of hostility. Black-danger stereotypes can also influence relatively automatic judgments. For example, priming participants with Black faces facilitates reactions to guns, but not tools (Payne, 2001). Also, recent neuroimaging data indicate greater activation of the amygdala, a brain structure involved in fear response, after brief (30 ms) exposure to Black compared to White faces (Cunningham et al., 2004).

Although it is likely that stereotypic and prejudicial associations toward Blacks are all interrelated to some extent, different types of associations may have discriminable effects on behavior. For example, Amodio and Devine (2006) demonstrated that implicit prejudice predicts consummatory behavior, whereas implicit stereotypes predict judgments and impression formation. Similarly, because rapidly detecting and responding to threats holds great importance for survival, danger-relevant stereotypes may be somewhat functionally distinct from danger-irrelevant stereotypes and from general prejudice. Rather than simply reflecting a general dislike for Blacks or a general accessibility of Black stereotypes, we hypothesize that danger stereotypes have distinct implications for behavior related to vigilance or threat detection. In this work, we focus on the well-studied link between threat and attention.

Threat and attention

Threatening stimuli such as snakes and spiders bias visual attention. Searching for a threatening stimulus in an array of non-threatening stimuli is faster and easier than searching for a non-

[☆] Primary support for this work was provided by National Science Foundation Continuing Grant 0642580 to the second and third authors. We would like to thank Stephanie Echols, Arina Goyle, Pamela Kaye, Debbie S. Ma, and Jessica Sim for their helpful comments.

^{*} Corresponding author. Fax: +1 773 834 9793.

E-mail address: nicoled@uchicago.edu (N.C. Donders).

threatening stimulus in an array of threatening stimuli (Öhman, Flykt, & Esteves, 2001). As many researchers have suggested (e.g. Fox, Russo, Bowles, & Dutton, 2001; Mogg & Bradley, 1998; Öhman, Flykt, & Lundqvist, 2000; Yiend & Mathews, 2001; Koster, Crombez, Verschuere, & DeHouwer, 2004), there are at least two distinct ways in which a stimulus can elicit attentional bias. Threatening stimuli may capture attention faster or hold attention longer than nonthreatening stimuli. These two processes may also act in tandem, such that threatening stimuli both capture and hold attention.

Posner (1980), see also Fox et al., (2001) devised a dot-probe task capable of separately assessing attentional capture and holding. Recently, Koster et al. (2004) used this attentional cueing task to show that fear-conditioned stimuli both capture attention faster and hold attention longer than neutral stimuli. In the preliminary fear-conditioning phase of this study, subjects were exposed to two different colored rectangles, one of which was consistently paired with an aversive noise burst designed to evoke threat. Koster et al. then measured attention for the fear-conditioned versus the neutral rectangle cue. Each trial of the task progressed as follows: first, a cue (either the fear-conditioned or the neutral rectangle) appeared on one side of a computer screen. The cue then disappeared, and a target (a square dot) appeared either in the same location as the cue had been or on the opposite side of the screen. Participants pressed one of two keys to indicate the side of the screen on which the target appeared. Reaction times (RTs) were the critical dependent variable.

The attentional cueing task includes two critical trial types: valid and invalid. On valid trials, the target appears in the same location as the cue. Because the cue grabs attention, orienting participants to the true location of the target, valid trials should facilitate detection of the target. Critically, if threatening stimuli capture attention faster than nonthreatening stimuli, valid trials on which the fear-conditioned cue appears should yield faster RTs than valid trials on which the neutral cue appears.

On invalid trials, the target appears in the location opposite the cue. Accordingly, participants must disengage from the cue in order to detect the target. Because the invalid cue holds attention, it should delay response to the target. If threatening stimuli hold attention longer than nonthreatening stimuli, invalid trials on which the fear-conditioned cue appears should yield slower RTs than invalid trials on which the neutral cue appears.

Compared to trials on which the neutral cue appeared, Koster et al. (2004) found that participants were faster to detect the target after a valid fear-conditioned cue, but slower to detect the target after an invalid fear-conditioned cue, indicating that threatening stimuli both capture attention more quickly and hold attention longer than neutral stimuli. We suggest that race has similar implications for attention. In fact, Levin (1996, 2000) has shown that searching for a Black face in an array of White faces is faster and easier than searching for a White face in an array of Black faces, suggesting that Black faces “pop out”. We propose that danger stereotypes should moderate this kind of effect. For individuals who do not link Blacks with danger, Blacks and Whites should not differ in their capacity to capture and hold attention.

An alternative to our position is that danger associations do not have discriminable effects on behavior, but instead simply reflect either a general dislike for Blacks or a general accessibility of all Black stereotypes. If this is true, prejudice toward Blacks or the accessibility of danger-irrelevant stereotypes may predict attentional bias. In contrast, we hypothesized that danger stereotypes would predict attentional allocation even when controlling for the effects of danger-irrelevant stereotypes and general prejudice.

Methods

Participants and design

Twenty-eight White University of Chicago undergraduates participated in exchange for money or course credit. Three participants were excluded from the analysis. For 2 participants, technical difficulties prevented full data collection. One additional participant was nonnative to the USA. Re-analysis including this participant yields no difference in the pattern of results. Overall, this left 25 participants (11 female, average age of 20.04).

The study involved two distinct tasks. First, the attentional cueing task measured differences in attentional capture and holding for Black, compared to White faces. This task employed a 2×2 within-subjects design, with face race (Black vs. White) and cue validity (valid vs. invalid) as repeated factors. The second task measured automatic stereotyping and prejudice. This task followed a 3 (association dimension: danger stereotype, danger-irrelevant stereotype, prejudice) \times 2 (cultural association: Black vs. White) within-subjects design.

Attentional cueing task

Our version of Koster et al. (2004)'s attentional cueing task measured attentional capture and holding by Black and White faces. Specifically, the task involved the following sequence of events (see Fig. 1): First, a fixation cross appeared in the center of the computer screen for 500 ms. Next, a face appeared either on the left or the right side of the screen for 40 ms.¹ The face was that of either a Black or a White male. The face then disappeared and was masked with scrambled pixels for 150 ms. A target then appeared on either the left or the right side of the screen for 1500 ms. This target, a square dot, was the average color of the Black and White faces used in the task. Participants responded to the target's location by pressing the “a” key if the dot was on the left, or the “l” key if the dot was on the right. If a timeout (response latency greater than 5000 ms) or an incorrect response occurred, three red X's appeared in the middle of the screen, accompanied by a buzzing sound. The intertrial interval was 1000 ms.

The task included two primary trial types: valid and invalid. Valid trials occurred when the dot appeared on the same side of the screen as the face. If Black faces capture attention faster than White faces, subjects should respond to the dot especially quickly on valid trials following a Black face rather than a White face. Invalid trials occurred when the dot appeared on the opposite side of the screen as the face. If Black faces hold attention longer than White faces, subjects should respond more slowly on invalid trials following a Black face rather than a White face.

Following Koster et al. (2004), *catch trials* and *digit trials* were included to ensure that participants followed instructions during the task. Catch trials ensured that participants responded to the dot instead of to the face. On catch trials, a face appeared, but no dot followed. Participants were instructed to make no response on these trials. On average, participants responded in error to a relatively small percentage (5.50%) of catch trials. Digit trials ensured that participants attended to the fixation cross at the beginning of each trial. On these trials, the fixation cross was replaced by a single-digit number presented for 100 ms, after which no face or dot followed. Participants typed in the digit that they saw. Participants were instructed to guess until they responded correctly. Only after responding correctly could they proceed to the next trial. On average, participants' first response was incorrect on only 2.29% of digit

¹ This short duration was chosen because we were interested in relatively automatic, rather than deliberative, face processing (see Cunningham et al., 2004).

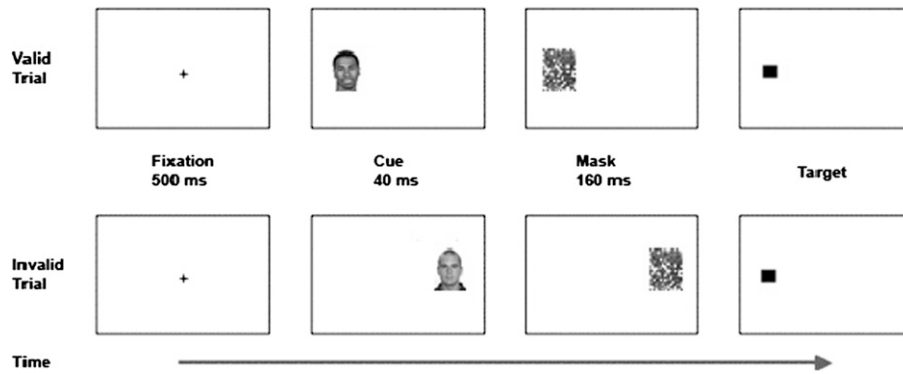


Fig. 1. Representation of a typical trial in the attentional cueing task.

trials, indicating that they oriented to the fixation cross at the beginning of most trials.

The task consisted of three blocks, with the first block composed of 42 practice trials. Block 2 served as a “buffer” block. Although participants thought that the practice phase had ended, Block 2 actually consisted of an additional 41 unanalyzed practice trials, allowing for an uninterrupted transition from practice to test trials. Throughout these two blocks, participants encountered 36 valid, 36 invalid, 6 catch, and 5 digit trials. Block 2 transitioned seamlessly into Block 3, the “test” block, for which RTs were analyzed. This block included 188 trials: 72 valid, 72 invalid, 24 catch, and 20 digit trials. In all blocks, the race of the face and the dot’s location were counterbalanced across trial type. Within each block, trials appeared in random order.

EAST

We measured automatic stereotyping and prejudice using the Extrinsic Affective Simon Task (EAST; De Houwer, 2003), a reaction time measure similar to the Implicit Association Test (Greenwald et al., 1998). The EAST presents a single stimulus on each test trial. In our modified version, some trials presented a face, which participants categorized as *Black* or *White* by pressing the “a” or “l” key, respectively. On other trials, words were presented in a blue-green typeface. Sometimes the word was slightly more blue than green, and sometimes it was slightly more green than blue. Participants classified each word as either *mostly blue* or *mostly green* using the same keys, “a” and “l”, respectively. The basic rationale underlying the EAST is that (1) participants tend to automatically process the word’s meaning even though it is irrelevant to the font color judgment, and (2) that by virtue of the face judgment task, the “a” and “l” keys acquire an extrinsic association with Black and White faces, respectively. As a result, when participants observe a word like *basketball*, those who associate the term with Blacks should be faster to classify the word if it appears in blue rather than green, because a blue word requires the same response as a Black face (“a” key) whereas a green word requires the same response as a White face (“l” key).

The main benefit of the EAST compared to other measures of automatic associations is that it allows separate indices of stereotyping and prejudice to be measured in a single task.² In our ver-

sion, the classified words belonged to several different categories (see Table 1). We used RTs from each word category to calculate an index of the association strengths between Black (relative to White) faces and (1) danger-relevant terms, (2) stereotypic terms unrelated to danger, and (3) positive and negative terms (for a description of calculations, see the Results section).

The EAST began with three practice blocks. First, participants practiced the race-categorization task (20 trials). Second, participants practiced the color-categorization task (32 trials). Third, they practiced the task with a random mix of race and word-color categorizations (50 trials). Finally, they performed a test block in which, again, the two trial types were randomly interspersed (320 trials). On each trial, participants had unlimited time to respond, but the trial did not terminate until the participant responded correctly.

Both the attentional cueing task and the EAST were programmed using PsyScope (Cohen, MacWhinney, Flatt, & Provost, 1993) and administered on Macintosh computers. Faces for the attentional cueing task were selected from the NimStim Face Stimulus Set (Tottenham et al., in press). Faces for the EAST were obtained from Ito and Urland (2003).

Procedure

The study was described as an investigation of vigilance. Participants completed the attentional cueing task and then the EAST, followed by a demographic questionnaire. The study lasted 45 min. Afterwards, participants were compensated with \$8 or course credit, thanked, and debriefed. No participant expressed suspicion about the purpose of the study.

Results

Mean level effects

Attentional allocation

We excluded data from trials on which participants responded either more quickly than 150 ms or more slowly than 1000 ms. Latencies from correct responses were submitted to a natural log transformation to reduce skew. Using data from only the attentional cueing task, we assessed whether on average, Black faces captured attention faster and held attention longer than White faces. We measured *attentional capture* by looking only at valid trials, where the face and the dot appeared on the same side of the screen. To the extent that Black faces capture attention faster than White faces responses to the dot should be faster when preceded by a Black face than a White face. Thus, attentional capture was calculated by subtracting RTs on trials where a Black face was presented (RT_{B_val} ; $M = 5.86$) from RTs on trials where a White face was presented

² De Houwer and De Bruyker (2007) recently questioned the EAST’s reliability and predictive validity. It is worth noting that our data offer some evidence of the EAST’s reliability and predictive validity as a measure of racial attitudes. That is, in addition to the primary prejudice index, described above, the EAST provides a second, separate index reflecting the extent to which Black faces facilitate responses to negative stereotypic words (e.g., “poor” and “boring”) and White faces facilitate positive stereotypic words (“athletic” and “smart”). In a pretest study, these separate estimates of prejudice generated by the EAST showed moderate intercorrelations ($r(77) = 0.37, p < .001$).

Table 1
Words used in the EAST

Association dimension		Predicted cultural association (stereotypic or prejudicial)	
		Black	White
Danger stereotypes		Danger, crime, violent, murder	Gentle, trust, peaceful, safety
Danger-irrelevant stereotypes	–/+	Poor, dishonest, welfare, lazy, athletic, jazz, rhythm, relaxed	Stuffy, boring, greed, selfish, math, success, smart, wealthy
Prejudice		Poison, awful, cancer, rotten	Wonderful, happy, delight, good

(RT_{W_val} ; $M = 5.88$). On average, Black faces did not capture attention more quickly than White faces, $M = 0.016$, $t(24) = 1.33$, $p = .20$.

Similarly, we calculated *attentional holding* by looking only at invalid trials, where the face and the dot appeared on opposite sides. If responses to the dot were slower when a Black face appears than when a White face appears, it would indicate that Black faces hold attention longer than White faces. Attentional holding was calculated by subtracting RTs on trials where a White face was presented (RT_{W_inv} ; $M = 5.87$) from RTs on trials where a Black face was presented (RT_{B_inv} ; $M = 5.86$). Again, there was no evidence in the current data that, on average, Black faces held attention longer than White faces, $M = 0.010$, $t(24) = 0.92$, $p = .37$.

An integrated attentional bias score reflecting both capture and holding was obtained by calculating the cue validity \times race interaction: $(RT_{W_val} - RT_{B_val}) + (RT_{B_inv} - RT_{W_inv})$. Higher values of integrated attentional bias reflect both faster capture and longer holding by Black faces compared to White faces. No significant validity \times race interaction was found, $t(24) = 1.44$, $p = .17$. Though trending in the predicted direction, our data provide no reliable evidence that Black and White faces affect attentional allocation differently on average.

Stereotype accessibility

We excluded data from trials on which participants responded either more quickly than 300 ms or more slowly than 3000 ms. Latencies from correct responses were submitted to a natural log transformation to reduce skew. We presented participants with six word categories that, based on cultural stereotypes and/or prejudice, should be associated with either Blacks or Whites (danger, safety, Black stereotypic, White stereotypic, nonstereotypic positive, nonstereotypic negative). We computed two averages (one for each response key: blue vs. green) for each of the six target-word categories. We used RTs for these categories to calculate indices of automatic accessibility for (1) danger stereotypes, (2) danger-irrelevant stereotypes, and (3) prejudice (i.e., negatively valenced stereotype-unrelated words; see Table 1).

To simplify description of the calculation of these indices, we focus on the danger words, since this category is the most theoretically relevant to our discussion. Our danger stereotype index reflects the extent to which participants associate Blacks (more than Whites) with danger, and Whites (more than Blacks) with safety. Mathematically, this is the interaction of face race \times predicted association: $(RT_{W_danger} - RT_{B_danger}) + (RT_{B_safety} - RT_{W_safety})$. Similarly, we used RTs for the Black and White danger-irrelevant stereotypic words to compute an index of danger-irrelevant stereotypes, and we used RTs for the negative and positive nonstereotypes to compute an index of prejudice.

On average, no significant race \times predicted association interaction emerged for danger stereotypes ($M = -0.004$, $t(24) = -0.09$, $p = .93$), danger-irrelevant stereotypes ($M = -0.04$, $t(24) = -0.72$, $p = .48$), or prejudice ($M = 0.027$, $t(24) = 0.81$, $p < .42$). While the accessibility indices showed no main effects, our primary interest in the current research concerned individual differences in danger stereotypes and their relation to attention.³ We therefore con-

ducted several regression analyses to assess whether danger stereotype accessibility predicts biased attentional allocation, controlling for danger-irrelevant associations.

Multiple regressions

We used multiple regressions to test our three indices of attention (capture, holding, and integrated bias) as a function of the EAST measures. For each regression we used a similar model, regressing each index of attentional bias on all three EAST indices (danger stereotypes, danger-irrelevant stereotypes, and prejudice) (see Footnote 3).

Overall attentional bias

We first analyzed whether danger stereotypes predict the integrated index of both attentional capture and holding. Controlling for danger-irrelevant stereotypes and general prejudice, EAST danger scores significantly predicted integrated attentional bias, $b = 0.25$, $t(21) = 3.14$, $p = .05$. Neither danger-irrelevant stereotype ($b = -0.09$, $t(21) = -1.62$, $p = .12$) nor prejudice ($b = 0.04$, $t(21) = 0.35$, $p = .73$) EAST scores significantly predicted biased attentional allocation. Thus, danger stereotypes, but no other associations that we measured, predicted differences in attentional capture and holding for Black versus White faces. As danger stereotype accessibility increases, Black faces capture attention faster and hold attention longer than White faces.

Of course, this integrated attentional bias index combines bias in attentional capture and holding. Next, we break down the attentional bias function into its two components, examining attentional capture and holding separately.

Attentional capture

To investigate whether attentional capture differed as a function of the EAST association strengths, we ran a multiple regression using the model described above. EAST danger scores significantly predicted attentional capture over and above the other EAST indices, $b = 0.15$, $t(21) = 2.62$, $p = .02$. Again, neither danger-irrelevant stereotype ($b = -.04$, $t(21) = -0.87$, $p = .40$) nor prejudice ($b = 0.003$, $t(21) = 0.04$, $p = .97$) EAST scores predicted attentional capture. This suggests that as participants show a stronger association between Blacks and danger, they orient more quickly to a Black face than a White face.

Faster attentional capture by Black versus White faces can be explained by Black faces capturing attention more quickly or by White faces capturing attention more slowly for individuals with highly accessible danger stereotypes. To distinguish between these two possibilities, we regressed RTs for valid trials on which a Black face appeared (B_val) on the three EAST indices. We repeated this process for White-face trials (W_val). On B_val trials, EAST danger scores significantly predicted the speed with which Black faces captured attention ($b = -0.27$, $t(21) = -2.79$, $p = .01$). Participants with highly accessible danger stereotypes showed faster attentional capture for Black faces. As before, neither danger-irrelevant stereotype ($b = 0.04$, $t(21) = 0.65$, $p = .52$), nor prejudice ($b = 0.07$, $t(21) = 0.57$, $p = .57$) EAST scores predicted the speed with which Black faces captured attention.

³ Inclusion of the secondary prejudice index and/or an index of stereotypic prejudice (Wittenbrink et al., 1997) had no effect on any of our regression models.

In contrast, on W_val trials danger stereotypes did not significantly predict the speed with which White faces captured attention ($b = -0.11$, $t(21) = -0.86$, $p = .40$) – nor did danger-irrelevant stereotypes ($b = 0.01$, $t(21) = 0.09$, $p = .93$), or prejudice ($b = 0.07$, $t(21) = 0.44$, $p = .66$) scores. In other words, danger stereotypes increase the speed with which participants orient to Black faces, without slowing attentional orienting toward White faces. Much like fear of spiders should increase the speed with which an individual orients to a tarantula relative to a tulip, danger stereotypes seem to increase the speed with which participants orient to a Black face, relative to a White face.

We can shed further light on these data by examining participants at low and high levels of danger stereotype accessibility. For individuals with weak danger stereotypes (EAST scores 1 SD below the mean), RTs on B_val and on W_val trials did not differ ($M = -0.02$, $t(21) = 0.95$, $p = .36$). However, for individuals with strong danger stereotypes (EAST scores 1 SD above the mean), RTs on B_val trials were significantly faster than RTs on W_val trials ($M = 0.05$, $t(21) = 2.91$, $p = .01$) (see Fig. 1). That is, individuals with strong danger stereotypes show faster attentional capture by Black than by White faces (see Fig. 2).

Attentional holding

Looking only at invalid trials, attentional holding reflects the extent to which responses are slower following a Black face compared to a White face. Here, danger stereotypes marginally predicted racial bias in attentional holding ($b = 0.1$, $t(24) = 1.88$, $p = .08$). Neither danger-irrelevant stereotype ($b = -0.06$, $t(24) = -1.52$, $p = .14$), nor prejudice ($b = 0.03$, $t(24) = 0.49$, $p = .62$) EAST scores predicted this measure. Thus, the marginally significant effect of racial associations on attentional holding was specific to danger stereotypes.

As in our attentional capture analysis, we looked separately at attentional holding by Black faces and by White faces, regressing RTs for only Black-face (B_inv) trials, and then for only White-face (W_inv) trials, on the three EAST indices. Neither danger stereotypes ($b = 0.01$, $t(24) = 0.06$, $p = .95$), danger-irrelevant stereotypes ($b = -0.04$, $t(24) = -0.38$, $p = .71$), nor prejudice ($b = 0.17$, $t(24) = 0.88$, $p = .39$) predicted attentional holding by Black faces. The same was true for White faces: neither danger stereotypes ($b = -0.09$, $t(24) = -0.60$, $p = .55$), danger-irrelevant stereotypes ($b = 0.015$, $t(24) = 0.14$, $p = .89$), nor prejudice ($b = 0.136$, $t(24) = 0.73$, $p = .47$) predicted attentional holding. It should be noted that these effects may be obscured because the dependent variable represents a simple average rather than a difference score; individual differences in mean RT therefore contribute to the error term, weakening the power to detect effects.




Fig. 2. Attentional capture for Black and White faces as a function of danger stereotype accessibility.

Participants with low danger stereotype accessibility (1 SD below the mean) showed no evidence of racially biased attentional holding. RTs on B_inv and W_inv trials did not differ ($M = -0.03$, $t(24) = -0.73$, $p = .48$). However, for participants with highly accessible danger stereotypes (1 SD above the mean), Black faces held attention marginally longer than White faces ($M = 0.03$, $t(24) = 1.88$, $p = .06$) (see Fig. 3).

Discussion

Research by Levin (1996, 2000) and Öhman et al. (2001) indicates that Black faces and threat-relevant stimuli may bias attention in similar ways. However, to our knowledge, the current data represent the first direct evidence that the extent to which Black faces bias attention is related to their perceived threat value. Race may bias attention in ways similar to fear-conditioned stimuli (see Koster et al., 2004). Danger stereotypes predict an increase in attentional capture, and a marginal increase in attentional holding by Black (rather than White) faces. These effects remain significant while controlling for the effects of danger-irrelevant stereotypes and prejudice. Further, neither danger-irrelevant stereotypes nor prejudice significantly predicts racial bias in attention. Rather, threat specifically relates to bias in attentional allocation.

It is interesting and perhaps instructive to note that other research has found a pronounced pattern of ingroup bias in attention. In particular, Eberhardt and colleagues (Eberhardt, Goff, Purdie, & Davies, 2004, Studies 2–4) have shown that predominantly White participants spend a disproportionate amount of time looking at White faces. By contrast, we observed only an interaction between attentional bias and danger stereotypes, with no overall bias to attend either to the ingroup or the outgroup. This divergence in findings may stem from differences in methodology. Eberhardt presented faces for a relatively long duration (450 ms), which may have allowed participants to engage in deeper and more individuated face processing, especially for the ingroup (Ito & Urland, 2003, 2005). By contrast, we presented faces quite briefly (40 ms), with the hope that this short presentation would engender quick threat detection without allowing for controlled regulation of the fear response (Cunningham et al., 2004). Indeed, we previously conducted a similar study, which presented faces for a longer duration (200 ms). Given this slight modification, White participants showed enhanced attention to the White faces: on average, White faces grabbed attention more quickly than Black faces, $p < .03$. Critically, however, we observed a relationship between danger stereotypes and attentional bias similar to the results reported above. Despite the average participant's bias to attend to White faces, those with stronger danger stereotypes oriented more quickly to the Black faces, $p < .03$. Though we do not




Fig. 3. Attentional holding for Black and White faces as a function of danger stereotype accessibility.

wish to make any strong assertions here, these results may reflect two additive processes: quick orientation to potential threats in the environment, and slightly slower, more individuated processing of ingroup members (Correll, Ito, & Urland, 2006; Ito & Urland, 2005).

Our findings suggest that different information associated with a social group can have distinct implications for behavior. This is in line with Amodio and Devine (2006), who have posited that stereotyping and prejudice take place via distinct cognitive and affective processing systems. Amodio and Devine argue that associating a Black target with the trait “poverty” and evaluating a Black target as “awful” reflect fundamentally different processes, each with distinct implications for behavior (see Wittenbrink, 2007). Regarding this distinction, danger stereotypes are an especially interesting case. On one hand, these threat-relevant associates are semantically specific, as are other, danger-irrelevant stereotypes such as “lazy” and “athletic”. However, threat-relevant associates are emotionally laden and critical for survival in ways that other types of semantic associates are not. Since danger stereotypes appear to have distinct implications for attention, threat-relevant stereotypes may represent a partially distinct category of information with separable implications for behavior.

Many factors may elicit differences in threat experienced upon exposure to Blacks compared to Whites. Any greater experience of threat upon exposure to a Black versus a White social target may cause biased attention, regardless of the source of this threat (Stephan & Stephan, 2000). In line with this, Richeson and Trawalter (2007) have shown that individuals high in external motivation to control prejudice (EM) show faster attentional capture by Black versus White faces in a dot-probe paradigm. High-EM individuals desire to respond in an unprejudiced manner not because they believe prejudice is wrong, but because they wish to conform to egalitarian social norms. Interracial interactions may thus provoke anxiety for these individuals (e.g. Richeson & Shelton, 2003), heightening the perceived threat value of Black faces. The current data are entirely consistent with this conjecture. Moreover, they indicate that stereotypes linking Blacks with danger are another potential source of threat that can elicit racially biased attention.

For individuals familiar with stereotypes linking Blacks with danger, Black faces may essentially serve as a fear-conditioned stimulus. Just as a stimulus can acquire threat value through repeated pairing with an aversive noise burst (Koster et al., 2004; see also Brosch & Sharma, 2005; Flykt, 1999), Black targets may acquire threat value due to, for example, a preponderance of violent Black characters in popular culture. Further research is necessary to test the possibility that the fear response toward Blacks is indeed acquired, and does not reflect a biologically prepared response to outgroup members (see Öhman & Mineka, 2001; Öhman et al., 2000). Likewise, it would be informative to identify the situational constraints for the acquisition of threat-relevant associations in order to better understand how society promotes danger-relevant stereotypic associations, and how the formation of these associations may be prevented.

References

- Amodio, D. M., & Devine, P. G. (2006). Stereotyping and evaluation in implicit race bias: Evidence for independent constructs and unique effects on behavior. *Journal of Personality and Social Psychology*, *91*, 652–661.
- Bradley, B. P., Mogg, K., Falla, S. J., & Hamilton, L. R. (1998). Attentional bias for threatening facial expressions in anxiety: Manipulation of stimulus duration. *Cognition & Emotion*, *12*, 737–753.
- Brosch, T., & Sharma, D. (2005). The role of fear-relevant stimuli in visual search: A comparison of phylogenetic and ontogenetic stimuli. *Emotion*, *5*, 360–364.
- Cohen, J. D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments, and Computers*, *25*, 257–271.
- Correll, J., Ito, T. A., & Urland, G. R. (2006). Event-related potentials and the decision to shoot: The role of threat perception and cognitive control. *Journal of Experimental Social Psychology*, *42*, 120–128.
- Cunningham, W. A., Johnson, M. K., Raye, C. L., Gatenby, J. C., Gore, J. C., & Banaji, M. R. (2004). Separable neural components in the processing of Black and White faces. *Psychological Science*, *15*, 806–813.
- De Houwer, J. (2003). The extrinsic affective Simon task. *Experimental Psychology*, *50*, 77–85.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 5–18.
- Dovidio, J. F., & Gaertner, S. L. (1986). *Prejudice, discrimination, and racism*. San Diego: Academic Press.
- Duncan, B. L. (1976). Differential social perception and attribution of intergroup violence: Testing the lower limits of stereotyping of Blacks. *Journal of Personality and Social Psychology*, *34*, 590–598.
- Eberhardt, 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., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*, *69*, 1013–1027.
- Flykt, A. (1999). *A threat imminence approach to human fear responding: Direction of threat, aversive contexts, and electrodermal responses*. Uppsala University: Uppsala, Sweden.
- Fox, E., Russo, R., Bowles, R., & Dutton, K. (2001). Do threatening stimuli draw or hold visual attention in subclinical anxiety? *Journal of Experimental Psychology: General*, *130*, 681–700.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*, *74*, 1464–1480.
- 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.
- Ito, T. A., & Urland, G. R. (2005). The influence of processing objectives on the perception of faces: An ERP study of race and gender perception. *Cognitive, Affective & Behavioral Neuroscience*, *5*, 21–36.
- Koster, E. H., Crombez, G., Verschuere, B., & DeHouwer, J. (2004). Selective attention to threat in the dot probe paradigm: Differentiating vigilance and difficulty to disengage. *Behaviour Research and Therapy*, *42*, 1183–1192.
- Levin, D. T. (1996). Classifying faces by race: The structure of face categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 1364–1382.
- Levin, D. T. (2000). Race as a visual feature: Using visual search and perceptual discrimination tasks to understand face categories and the cross-race recognition deficit. *Journal of Experimental Psychology: General*, *129*, 559–574.
- Öhman, A., Flykt, A., & Esteves, F. (2001). Emotion drives attention: Detecting the snake in the grass. *Journal of Experimental Psychology: General*, *130*, 466–478.
- Öhman, A., Flykt, A., & Lundqvist, D. (2000). Unconscious emotion: Evolutionary perspectives, psychophysiological data and neuropsychological mechanisms. In R. D. Lane & L. Nadel (Eds.), *Cognitive neuroscience of emotion* (pp. 296–327). New York, NY: Oxford University Press.
- Öhman, A., & Mineka, S. (2001). Fears, phobias, and preparedness: Toward an evolved module of fear and fear learning. *Psychological Review*, *108*, 483–522.
- 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.
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*, 3–25.
- Richeson, J.A., & Trawalter, S. (2007). unpublished manuscript
- Richeson, J. A., & Shelton, J. N. (2003). When prejudice does not pay: Effects of interracial contact on executive function. *Psychological Science*, *14*, 287–290.
- Sagar, H. A., & Schofield, J. W. (1980). Racial and behavioral cues in Black and White children's perceptions of ambiguously aggressive acts. *Journal of Personality and Social Psychology*, *39*, 590–598.
- Stephan, W. G., & Stephan, C. W. (2000). An integrated threat theory of prejudice. In S. Oskamp (Ed.), *Reducing prejudice and discrimination* (pp. 23–45). Mahwah, NJ: Lawrence Erlbaum.
- Tottenham, N., Tanaka, J., Leon, A.C., McCarry, T., Nurse, M., Hare, T.A., et al. (in press). The NimStim set of facial expressions: Judgments from untrained research participants. *Psychiatry Research*.
- Wittenbrink, B. (2007). Measuring attitudes through priming. In B. Wittenbrink & N. Schwarz (Eds.), *Implicit measures of attitudes* (pp. 17–58). New York, NY: Guilford Press.
- Wittenbrink, B., Judd, C. M., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology*, *72*, 262–274.
- Yiend, J., & Mathews, A. (2001). Anxiety and attention to threatening pictures. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, *54a*, 665–681.