Journal of Experimental Social Psychology 55 (2014) 74–79

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

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

On the flexibility of attention to race ☆

Joshua Correll [a,*], Steffanie Guillermo [a], Julia Vogt [b]

[a] University of Colorado at Boulder, Boulder, CO, USA
[b] University of Reading, Reading, Berkshire, UK

HIGHLIGHTS

- By introducing a goal manipulation, the current study examines whether attention to racial outgroups is obligatory or flexible.
- In the absence of a goal, White participants showed clear evidence of preferential attention to Black faces.
- By contrast, in the presence of a goal that was unrelated to race, White participants showed no evidence of attention to Black faces.
- A goal induction eliminated attention to Black faces even for White participants with highly accessible stereotypes linking Blacks with the concept of danger.

ARTICLE INFO

Article history:
Received 12 November 2013
Revised 20 May 2014
Available online 28 May 2014

Keywords:
Attention
Race
Goals
Threat

ABSTRACT

Research on the flexibility of race-based processing offers divergent results. Some studies find that race affects processing in an obligatory fashion. Other studies suggest dramatic flexibility. The current study attempts to clarify this divergence by examining a process that may mediate flexibility in race-based processing: the engagement of visual attention. In this study, White participants completed an exogenous cuing task designed to measure attention to White and Black faces. Participants in the control condition showed a pronounced bias to attend to Black faces. Critically, participants in a goal condition were asked to process a feature of the stimulus that was unrelated to race. The induction of this goal eliminated differential attention to Black faces, suggesting that attentional engagement responds flexibly to top-down goals, rather than obligatorily to bottom-up racial cues.

Social cognitive research on race often involves paradigms in which participants passively view a race-related visual cue. For example, an evaluative priming task might present a Black or White face (which participants simply observe) followed by a word that participants classify as positive or negative. The effect of race involves the degree to which race affects subsequent judgments: relative to a White face, does the presentation of a Black face lead participants to classify negative words more quickly, but positive words more slowly?

In the current study, we examine the degree to which sensitivity to race is flexible (rather than involuntary). This question can be examined by modifying the standard paradigm: instead of allowing participants to passively observe a face, researchers might instruct them to attend to another aspect of the stimulus. Macrae, Bodenhausen, Milne, Thorn, and Castelli (1997) utilized such a paradigm to examine gender-based processing. They presented participants with photographs of male and female faces, some of which had a dot superimposed on them. Participants in a goal condition were asked to indicate the presence or absence of the dot. If gender influences behavior *even when participants are trying to attend to other information* (the dot), it suggests that the effects of gender are at least partially involuntary.[1] By contrast, if the goal manipulation eliminates effects of gender, it suggests greater flexibility: goals can override the influence of gender (in fact, this is what they found).

Researchers have adapted Macrae's paradigm to study race. Ito and Urland (2005; see also Ito & Urland, 2003) presented Black and White faces, some with dots, some without dots. Participants in one condition were asked to classify the faces by race; in another condition they were asked to detect the dot. The researchers found effects of race on event-related brain potential (ERP) components, including the P200, N200 and P300, which reflect processes related to early visual attention and working memory (Ito & Cacioppo, 2000; Luck & Hillyard, 1994). Instructing participants to attend to the dot reduced these effects, but – critically – it did not eliminate them, suggesting at least partially involuntary processing of race. In contrast, Wheeler and Fiske (2005), using essentially the same task, found evidence that race-based processing was more flexible. In a control condition, Black targets prompted increases in amygdala activity (Study 1) and stereotype activation

☆ This material is based on work supported by National Science Foundation Continuing Grant 0642580. We thank Mark Ohrtman for valuable contributions to this research.
* Corresponding author at: Department of Psychology & Neuroscience, University of Colorado at Boulder, Campus Box 345, Boulder, CO, 80309-0345, USA.
E-mail address: joshua.correll@colorado.edu (J. Correll).

[1] This kind of task may also implicate the efficiency with which participants process gender, though that phenomenon is less central to the current research question.

http://dx.doi.org/10.1016/j.jesp.2014.05.013
0022-1031/

(Study 2). However, when participants were asked to detect the dot, there was no evidence of increased amygdala activity or stereotype activation to Black targets. Contrary to Ito and Urland, these effects suggest greater flexibility of race-based processing in that it can be avoided when participants focus on other information (see also Kurzban, Tooby, & Cosmides, 2001).

Notably, Ito and Urland (2005) tested early stages of visual/cognitive processing (e.g., the N200 occurs roughly 190 ms after stimulus onset). By contrast, in Study 1, Wheeler and Fiske examined blood oxygen level-dependent responses that occur slowly over several seconds. In their second study, they imposed a two-second stimulus onset asynchrony (SOA) between the face and the lexical stimulus. One possibility, then, is that race-based processing is relatively inflexible in early stages of visual processing, but that top-down goals modulate the engagement of visual attention, allowing flexibility in subsequent processing. In the present research, we examine this question by testing whether goals influence visual attention.

Previous research has shown that White participants typically attend to Black faces, perhaps as a consequence of the stereotypic association between Blacks and danger. For instance, Donders, Correll, and Wittenbrink (2008) showed preferential attention to Black faces, especially for participants who associate Blacks with danger (see also Trawalter, Todd, Baird, & Richeson, 2008). This work extends a long tradition of research showing that people attend to threat-relevant stimuli. Snakes and spiders "pop out" in arrays of flowers and mushrooms (Öhman, Flykt, & Esteves, 2001; Öhman, Lundqvist, & Esteves, 2001), and participants detect a dot more quickly if it appears in the location of a snake than if it appears in the location of a nonthreatening image (Lipp & Derakshan, 2005). Öhman and his colleagues argued that, over the course of evolution, humans developed specialized cognitive and neural processes that detect threat quickly, efficiently, and – critically for the current investigation – involuntarily.[2] Racial outgroups may evoke this kind of processing either because they strike perceivers as foreign and different and therefore more likely to pose a threat (Olsson, Ebert, Banaji, & Phelps, 2005) or because phylogenetic threat-detection processes can be applied to stimuli that perceivers have ontogenetically learned to fear (e.g., guns). A number of recent studies suggest that perceived threat value, more than evolutionary significance, influences attention (Brosch & Sharma, 2005; Notebaert, Crombez, Van Damme, De Houwer, & Theeuwes, 2011; Young, Brown, & Ambady, 2012). These arguments suggest that, if race is associated with threat, it should capture attention even when participants are trying to focus attention elsewhere.

Other work suggests that goals can override attention to threat. Vogt, De Houwer, Crombez, and Van Damme (2013; see also Vogt, De Houwer, & Crombez, 2011) argue that, in situations where goals compete with threat cues for attention, goals win. They (found that participants directed spatial attention to nonthreatening but goal-relevant stimuli even when threatening stimuli were presented. They argue that goals can eliminate the (otherwise robust) tendency to orient to threat, suggesting that engagement of visual attention is determined by top-down processes, not evolved modules.

The present study tests whether attention can similarly "gate" the processing of race. Although Ito and Urland (2005) suggest that early visual processes obligatorily respond to race, goals may influence attentional processes that precede (and potentially initiate) the kind of effects Wheeler and Fiske observed (cf. Pourtois, Schettino, & Vuilleumier, 2013; Vuilleumier & Huang, 2009). Attention may thus intervene between obligatory forms of early visual processing – which are sensitive to race – and later processing that allows for greater top-down control (cf. Treisman, 1964; Wolfe & Horowitz, 2004).

[2] This process may be more accurately characterized as conditional automaticity in that it depends on the perceiver's capacity to visually process the stimulus. If the task profoundly taxes perceptual capacity, participants may not process the threat cue at all (Pessoa et al., 2005).

White participants were presented with images of White and Black faces in an exogenous cueing task (Fox, Russo, Bowles, & Dutton, 2001; Koster, Crombez, Verschuere, & De Houwer, 2004; Posner, 1980). Critically, participants in goal conditions were instructed to respond to a colored frame surrounding the face. In the control condition, the frames were irrelevant. If race influences attention in a relatively involuntary fashion, participants in the goal conditions should attend to Black faces in spite of the frame task. We refer to this as the *involuntary attention hypothesis* (H_1). However, if attention to race is more flexible, the goal induction should minimize – and perhaps eliminate – differential attention to Black faces. We refer to this as the *flexible attention hypothesis* (H_2). These alternatives operate as competing hypotheses. We presented faces for only 100 ms in order to investigate whether top-down processes intervene at early stages of attentional allocation (very soon after the seemingly obligatory response to race demonstrated by Ito & Urland, 2003). These brief presentations extend existing work on top-down influences (e.g.., Vogt et al., 2013) by examining extremely early attentional responses.

Methods

Participants

Participants were 48 White undergraduate students at the University of Chicago who participated for course credit. Computer malfunctions for three participants and one clear outlier (Cook's $d = 0.17$ and studentized deleted residual $= -3.54$) left 44 participants (23 female) for analysis.

Design and procedure

Each trial of the exogenous cueing task began with a fixation cross ("+") presented for 500 ms. Next, a face appeared on the left or right side of the screen for 100 ms. Following a 100 ms pause, a dot appeared on either the left or right side of the screen for 1500 ms. The location of the dot was randomized across trials. Participants were instructed to indicate the location of the dot (left or right) as quickly as possible, by pressing one of two buttons.

The cueing task employed two different types of trials: valid and invalid. On valid trials, the dot appeared on the same side of the screen as the face. On invalid trials, the dot appeared on the opposite side. If the face (a) captures attention on valid trials, drawing attention to the true location of the dot, and (b) holds attention on invalid trials, preventing participants from disengaging from the incorrect location, participants should identify the dot more quickly on valid trials than on invalid trials. The difference between latencies on invalid and valid trials thus served as our primary measure of attentional engagement. This difference was calculated separately for White and Black cues.

The task also included digit trials on which a number from 0 to 9 was presented instead of the fixation cross, and participants were instructed to type the number they saw (cf. Koster et al., 2004). These trials were utilized to determine whether participants fixated on the center of the screen at the beginning of each trial (participants correctly identified the digits on more than 90% of trials). A 42-trial practice phase (30 valid, 10 invalid and 2 digit) was followed by a 172-trial test phase (120 valid, 40 invalid and 12). Accordingly, the ratio of valid to invalid trials was 3:1. These trial frequencies were based closely on previous work with the exogenous cueing task (Koster et al., 2004). Stimuli included photographs depicting the faces of five White and five Black males. Two versions of each photograph were prepared, one with a green frame, one with a blue frame. Race, frame color and cue location (left or right) were randomized across trials.

Participants were randomly assigned to one of three goal conditions: find-blue-frames, find-green-frames, or a no-goal control condition. After identifying the location of the dot, participants in the goal conditions were instructed to press the space bar if they had seen a blue (or green) frame surrounding the picture of the face.[3] Participants were paid 10 cents for every correct frame detection and lost 10 cents for every false alarm. In the control condition, the frames were not mentioned. The research question involves a distinction between (a) participants with no goal and (b) participants with a goal to attend to *either* frame (blue or green). We included two colors simply to minimize the possibility of color-specific effects. The important point is that, in both goal conditions, participants have a goal that is not related to race. To maximize our power to detect differences between the no-goal condition and the two goal conditions (combined), random assignment over-weighted the no-goal condition relative to the find-blue and find-green conditions (n's = 19, 11, and 14, respectively).

After the cueing task, participants completed a modified version of the Extrinsic Affective Simon Task (EAST; De Houwer, 2003).[4] This task was identical to that used by Donders et al. (2008) and served to assess stereotypes linking Black men with the concept of danger. To simplify the findings, we present the primary results without reference to the EAST (which did not moderate the critical effects). But we address the measure later in the Results section to clarify our interpretation.

The study employed a 3 (goal condition: find-green, no-goal, find-blue) × 2 (face race: Black, White) × 2 (frame color: blue, green) × 2 (cue validity: valid, invalid) mixed-model design with goal condition varying between participants and all other factors varying within.

Results

We excluded test trials on which participants responded incorrectly (3.47%) or on which they responded faster than 100 ms (2.17%) or slower than 1000 ms (1.11%). Because the latencies were positively skewed (skewness = 1.45), they were log transformed (resulting in skewness = 0.26). We computed averages for each cell of the face race × frame color × cue validity design, yielding eight averages per participant.

We submitted these means to a 3 (goal condition) × 2 (face race) × 2 (frame color) × 2 (cue validity) mixed-model analysis, in which goal condition varied between participants and all other factors varied within. To test the predicted effects of goal condition, we specified linear and quadratic orthogonal contrasts (linear: find-green = −1, no-goal = 0, find-blue = +1; quadratic: find-green = −1, no-goal = +2, find-blue = −1).

It is important to note that interactions between validity and either race or frame color indicate differential attention. For example, the frame color × validity interaction might emerge because, relative to green frames, participants respond more quickly to blue frames on valid trials (suggesting attentional capture) but more slowly to blue frames on invalid trials (suggesting attentional holding). On average across goal condition, the color × validity interaction was marginally significant, $F(1,41) = 3.61$, $\eta_p^2 = .081$, $p < 0.065$, suggesting that participants attend slightly more to blue than green on average. More critically, the race × validity interaction was significant, $F(1,41) = 4.16$, $\eta_p^2 = .092$, $p < .048$, suggesting preferential attention to Black faces over White faces.

The analysis yielded a three-way interaction between the linear goal condition contrast, frame color and validity, $F(1,41) = 66.09$, $\eta_p^2 = .617$, $p < .0001$ (the quadratic goal × frame × validity interaction was not significant, $F(1,41) = 2.22$, $\eta_p^2 = .051$, $p < .15$) (see Table 1 and Fig. 1 for means). To interpret this effect, we examined the simple frame ×validity interaction in each goal condition. For participants who were paid to find green, the frame × validity interaction, $F(1,13) = 11.54$, $\eta_p^2 = .470$, $p < .005$, indicated greater attention (faster capture, longer holding) to green frames. For participants who were paid to find blue, the frame × validity interaction, $F(1,10) = 38.82$, $\eta_p^2 = .795$, $p < 0001$, was opposite in direction, indicating greater attention to blue frames. For no-goal participants, the interaction was not significant, $F(1,18) = 0.01$, $\eta_p^2 = .001$, $p < .95$, indicating no difference in attention to frames of different colors. We also decomposed the interaction by separately examining valid and invalid trials. To assess attentional capture, we modeled responses to valid trials as a function of goal condition and frame color. The linear goal × frame interaction was significant, $F(1,41) = 10.98$, $\eta_p^2 = .211$, $p < .002$, indicating the same pattern of linearly increasing attention to blue frames (the corresponding quadratic interaction was not significant, $F(1,41) = 1.21$, $\eta_p^2 = .029$, $p < .28$). To assess holding, we modeled RTs on only the invalid trials as a function of goal condition and frame color. Again, the linear goal × frame interaction emerged, $F(1,41) = 38.99$, $\eta_p^2 = .487$, $p < .0001$ (the quadratic interaction was not significant, $F(1,41) = 0.54$, $\eta_p^2 = .013$, $p < .47$). Simply put, paying participants to detect either blue or green frames prompted them to attend to the goal-relevant color. These effects are not surprising – they simply show that people attend to stimuli for which they are rewarded – but they are important because they suggest that we effectively manipulated participants' goals.

The analysis also revealed a three-way *quadratic* goal × face race× validity interaction, $F(1,41) = 5.37$, $\eta_p^2 = .116$, $p < .026$ (the linear goal × face race× validity interaction was not significant, $F(1,41) = 0.00$, $\eta_p^2 = .000$, $p < .99$). The quadratic contrast compares the no-goal condition with the average of the two goal conditions, thus it assesses the difference between having no goal and having some goal. The three-way interaction reflects the degree to which either goal reduces attention to Black faces. To interpret this effect, we again tested the simple race × validity interaction in each goal condition. For participants in the no-goal condition, the race × validity interaction was significant, $F(1,18) = 11.47$, $\eta_p^2 = .389$, $p < .004$, indicating preferential attention to Black faces (see Table 1 and Fig. 2). However, a goal eliminated this effect: the race × validity interaction was not significant among participants in either the find-green, $F(1,13) = 0.07$, $\eta_p^2 = .005$, $p < .80$, or find-blue goal condition, $F(1,10) = 0.09$, $\eta_p^2 = .009$, $p < .77$ (nor was it significant when we pooled the two goal conditions for a more powerful test: $F(1,24) = 0.17$, $\eta_p^2 = .007$, $p < .69$).

We also decomposed the three-way interaction by separately analyzing valid and invalid trials. The test of the valid trials provides perhaps the most rigorous test of the primary research question. Because participants generally respond quickly on valid trials, goals would have to influence processing at very early stages to have any effect on latencies. We analyzed valid trials as a function of goal condition and race. The main effect of race was not significant, $F(1,41) = 1.91$, $\eta_p^2 = .045$, $p < .18$, offering no evidence that Black faces captured attention in general. But, as predicted by the flexibility hypothesis, the quadratic goal × race interaction was marginally significant, $F(1,41) = 3.39$, $\eta_p^2 = .076$, $p < .073$ (the linear goal × race effect was not significant, $F(1,41) = 1.41$, $\eta_p^2 = .033$, $p < .25$). In the no-goal condition, Black faces captured attention more than Whites, $F(1,18) = 4.97$, $\eta_p^2 = .216$, $p < .039$, but a goal induction eliminated this effect. Black faces did not

[3] Importantly, because this task involves the simple detection of color, it should require minimal cognitive resources (Treisman, 1998), allowing sufficient resources for participants to process the stimuli (see Footnote 1; Pessoa et al., 2005).

[4] The EAST presents a single stimulus on each trial. On some trials, this stimulus is a male face (either Black or White), and on other trials it is a word (presented in either a blue or a green typeface). When a face appears, participants are asked to classify it by race, pressing the A key for a Black face and the L key for a White face. By virtue of this task, each key becomes associated with one racial group (A = Black; L = White). When a word appears, participants are asked to classify it by color, pressing the A key for a word printed in blue and the L key for a word printed in green. Each word (e.g., "criminal") is presented both in blue (requiring the button associated with Blacks) and in green (requiring the button associated with Whites). If a participant is faster to classify a word when it is printed in blue (rather than green), it suggests an association with Blacks. The EAST used a different set of faces, and presented a variety of words designed to measure both stereotyping and prejudice (for complete details, see Donders et al., 2008). Of particular interest, the task also involved a set of words clearly related to the dimension of danger—safety: danger, crime, violent, murder, gentle, trust, peaceful, safety. In three practice blocks, participants performed the race-categorization task (20 trials), the color-categorization task (32 trials), and a random mix of race and color categorizations (50 trials). Then, they performed a test block (320 trials) in which faces and words were randomly interspersed. On each trial, participants had unlimited time to respond, but the trial did not terminate until the participant responded correctly.

Table 1. Means and standard deviations of log-transformed response times by goal condition, race, frame color and validity (top panel) and means and standard deviations for attention indices (frame color and race) by goal condition (bottom panel).

Race	Frame	Validity		Goal condition					
				Find green		No goal		Find blue	
				M	*s*	*M*	*s*	*M*	*s*
Black	Blue	Valid	*a*	6.088	0.138	5.813	0.145	5.938	0.123
		Invalid	*b*	6.335	0.126	6.012	0.155	6.412	0.181
	Green	Valid	*c*	6.025	0.141	5.782	0.157	5.973	0.132
		Invalid	*d*	6.415	0.132	5.994	0.168	6.205	0.157
White	Blue	Valid	*e*	6.090	0.131	5.833	0.157	5.899	0.143
		Invalid	*f*	6.335	0.118	5.990	0.143	6.371	0.176
	Green	Valid	*g*	6.045	0.146	5.817	0.168	5.990	0.129
		Invalid	*h*	6.429	0.108	5.962	0.155	6.215	0.159

Index	Formula		Goal condition					
			Find green		No goal		Find blue	
			M	**s**	*M*	**s**	*M*	**s**
Blue grab	$(c+g)-(a+e)$	*w*	-0.108^{+}	0.233	-0.046^{**}	0.129	0.126^{*}	0.153
Blue hold	$(b+f)-(d+h)$	*x*	-0.175^{*}	0.231	0.046	0.165	0.362^{**}	0.262
Blue attn	$w+x$		-0.282^{**}	0.307	−0.001	0.140	0.488^{**}	0.257
Black grab	$(e+g)-(a+c)$	*y*	0.022	0.082	0.055^{*}	0.104	−0.021	0.064
Black hold	$(b+d)-(f+h)$	*z*	−0.014	0.094	0.054^{+}	0.149	0.031	0.098
Black attn	$y+z$		0.009	0.141	0.109^{**}	0.137	0.009	0.122

$^{+}$ $p < .15$.
* $p < .05$.
** $p < .01$.

differentially capture attention among participants in either the find-green, $F(1,13) = 1.14$, $\eta_p^2 = .081$, $p < .31$, or find-blue condition, $F(1,10) = 0.86$, $\eta_p^2 = .080$, $p < .38$ (pooling the two goal conditions: $F(1,24) = 0.10$, $\eta_p^2 = .004$, $p < .76$).

We also analyzed the invalid trials. On average, Black faces did not preferentially hold attention, $F(1,41) = 1.56$, $\eta_p^2 = .037$, $p < .22$, nor did goal condition moderate holding for Black faces, F's$(1,41) = 0.81$, 1.49, $\eta_p^2 = .019$, .035, p's < .38, .23 (for the quadratic and linear goal × race interactions, respectively). These results suggest that the effects on attention are primarily driven by attentional capture (a pattern that mirrors the findings of Donders et al., 2008). Analysis of capture and the integrated measure of attention both suggest that a goal reduces (even eliminates) preferential attention to Black faces.

It seems plausible that the effects of the goal induction might constrain attention only on goal-relevant trials (Vogt et al., 2013). Participants in the find-blue condition might inhibit attention to Black faces when the goal-relevant cue (a blue frame) is present. In the absence of that cue (a green frame), they might still attend to race. To examine this possibility, we conducted analyses on all participants in the goal conditions, examining attentional capture on both goal-present and goal-absent trials. Participants showed no evidence of preferential attention to Black faces in either case (goal-present: $M = -0.005$, $s = 0.077$, $t(24) = -0.33$, $\eta_p^2 = .005$, $p < 0.75$; goal-absent: $M = 0.010$, $s = 0.061$, $t(24) = 0.82$, $p < 0.42$), nor was there evidence that attention to Blacks was lower on goal-present trials relative to goal-absent trials ($M = 0.015$, $s = 0.113$, $t(24) = 0.67$, $\eta_p^2 = .018$, $p < 0.52$).

In summary, the data show that participants who were paid to identify green or blue frames paid more attention to the color for which they were rewarded. Critically, the goal of attending to *either* frame reduced preferential attention to Black faces, especially attentional capture. This pattern clearly supports the flexible goal hypothesis (H_2) and suggests that race does *not* involuntarily garner attention.

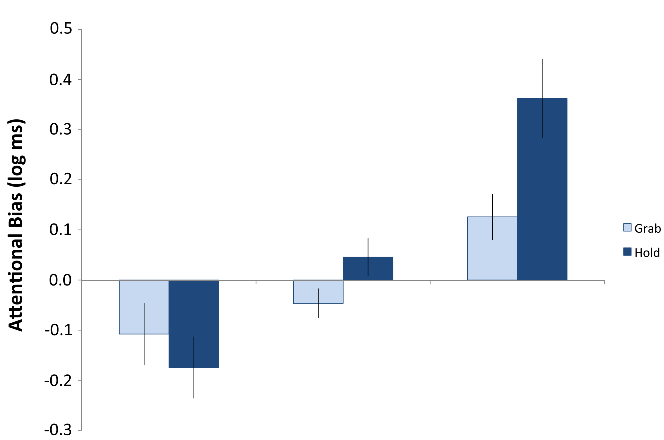

Fig. 1. Preferential attention to blue (rather than green) frames (means and standard errors). For valid trials, attentional capture (the "grab" index) was calculated as $blue_{grab} = (green_{valid} - blue_{valid})$. For invalid trials, attentional holding (the "hold" index) was calculated as $blue_{hold} = (blue_{invalid} - green_{invalid})$. Note that both indices collapse across Black and White faces.

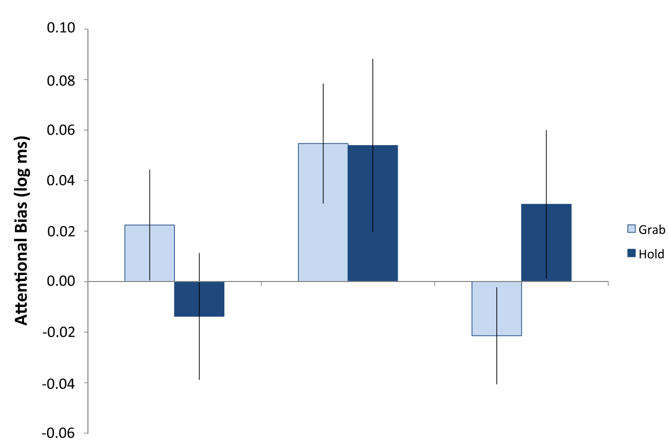

Fig. 2. Preferential attention to Black (rather than White) faces (means and standard errors). For valid trials, attentional capture (the "grab" index) was calculated as $Black_{grab} = (White_{valid} - Black_{valid})$. For invalid trials, attentional holding (the "hold" index) was calculated as $Black_{hold} = (Black_{invalid} - White_{invalid})$. Note that both indices collapse across blue and green frames.

Individual differences in stereotype accessibility

Donders et al. (2008) found that preferential attention to Black faces depended on stereotype accessibility. Participants who associated Blacks with danger demonstrated more pronounced attention to Black faces. Perhaps race involuntarily biases attention, but only for participants with highly accessible stereotypes about Blacks and danger. In other words, a Black face may constitute a meaningful threat (and involuntarily bias attention) only among participants who see the group as threatening in the first place. To examine this question, we computed an index of danger-based stereotyping for each participant, based on correct log-transformed response times during the EAST. Briefly, this index reflects the tendency to (a) respond to danger-related words more quickly when the response required participants to use a button associated with Blacks, rather than a button associated with Whites, and (b) respond to safety-related words more quickly when the response required the White rather than the Black button (for details, see Donders et al., 2008). Higher scores suggest a stronger stereotypic association between Blacks and danger. We reanalyzed the exogenous cueing task as a function of goal condition, race, cue validity, and danger stereotypes (continuously measured and mean centered).

To simplify our description of the results (which involve a four-way interaction), we computed an index of attention to Black faces, collapsing across frame color: $\text{Black}_{\text{attn}} = (\text{White}_{\text{valid}} - \text{Black}_{\text{valid}}) + (\text{Black}_{\text{invalid}} - \text{White}_{\text{invalid}})$. This index represents the race × validity interaction and constitutes the dependent variable in the following analysis. The primary analysis (described above) showed that a goal induction reduced attention to Blacks (i.e., the effect of the quadratic goal contrast on the race × validity interaction). Our present question involves the possibility that this reduction might be less pronounced for participants who strongly associate Blacks with danger, suggesting that (at least for participants with accessible stereotypes) Black faces do engage attention in an involuntary fashion. There was no evidence that danger stereotypes moderated the effects of goal condition, Fs$(1,38) = 0.02, 0.43$, $\eta^2_p = .001, .011$, $p < .88, .52$ (for the linear and quadratic interactions, respectively). And, even in this relatively complex model, the simple effect of quadratic goal condition (for participants at average levels of danger stereotyping) was significant, Fs$(1,38) = 5.24$, $\eta^2_p = .121$, $p < .028$. Controlling for danger stereotypes and their interactions with condition, the goal manipulation significantly reduced preferential attention to Black faces. It is also interesting to note that danger stereotypes were positively (though not significantly) correlated with capture to Black faces in the control condition, $r(17) = .36$, $p < .13$. In line with Donders et al. (2008), this suggests that threat stereotypes promote attention to Black faces. However, this relationship was non-significantly reversed in the goal conditions (combined to increase sample size), $r(23) = -.32$, $p < .12$, offering no evidence of a threat-attention relationship when goals were salient.

Discussion

This study investigated whether people involuntarily attend to race. In the absence of a competing task, participants showed preferential attention to Black faces, but a goal induction eliminated this effect. For participants who were asked to attend to frame color, Black faces neither captured attention faster nor held attention longer than White faces. This flexibility emerged even for participants with strong stereotypic associations and was more pronounced on indices assessing (relatively fast) attentional capture.

Our data address a potential intervening process between early visual processing and downstream stages related to evaluation and semantic association. Visual attention seems to respond flexibly to participants' goals. Latencies for valid trials (which assess attentional capture) were roughly 580 ms after the onset of the face (380 mean latency + 200 SOA), suggesting that a goal induction can minimize attention to race by that point in the processing stream. The time course of the effects reported here falls between the early stages of visual processing which show involuntary processing of race at 190 ms (Ito & Urland, 2005), and later stages of evaluative processing, which show goal-based flexibility between 2 and 6 seconds (Wheeler & Fiske, 2005). These findings indicate that flexibility in attention emerges early, raising the possibility that it operates as a transition point in the perception of race.

A number of details strengthen our confidence in these effects. First, it has been argued that, if a task is sufficiently demanding, it can preclude attention to *any* threat cue (Pessoa, McKenna, Gutierrez, & Ungerleider, 2002; Pessoa, Padmala, & Morland, 2005). It is therefore important that our goal involved the simple detection of color, an efficient process that requires minimal resources (Treisman, 1998). Perceptually, participants in the goal condition should have had ample capacity to process the faces we presented: if those faces *compelled* attention, participants should have attended to them. Despite this capacity, participants with the frame-detection goal showed no evidence of preferential attention to Black faces.

Further, previous work suggests that when participants view threatening stimuli, motor responses may be inhibited (Lang, Bradley, & Cuthbert, 1997). Participants may "freeze" upon presentation of a threatening stimulus, which distorts measures of both capture and holding by increasing reaction times. Mogg, Holmes, Garner, and Bradley (2008) suggested that the exogenous cueing task may thus systematically *underestimate* attentional capture and overestimate attentional holding. They suggest that the combined index of attention controls for any response slowing by integrating both capture (where freezing should weaken the effect) and holding (where freezing should inflate the effect). We found similar patterns whether we used the integrated index or the index of attentional capture. The fact that response slowing may weaken capture estimates suggests that these effects may be even stronger than they appear.

Recent work has shown that the allocation of attention is more flexible than often assumed. Vogt and her colleagues suggest a central role for top-down goals in the allocation of attention (e.g., Vogt, Lozo, Koster, & De Houwer, 2011; Vogt et al., 2013). Their work finds that attention is preferentially allocated to emotional stimuli only when those events are relevant to the observer's current goal or in the absence of competing goals. The present study advances our understanding in three ways. First, it extends the findings to the domain of race. Second, it shows that goals can reduce attention to race on all trials, even when goal-relevant cues were not present. Third, given the brief duration of the cues in this study (100 ms), the results suggest that goals modulate fairly early attention, which is often assumed to be purely stimulus-driven (e.g., Wolfe & Horowitz, 2004).

These findings suggest that attention to race is not obligatory. Although White participants typically preferentially attend to Black faces, a goal induction eliminates this bias. We observed this flexibility even though the primary task requires minimal cognitive resources, and even among participants who hold strong stereotypic associations between Blacks and danger. By examining attentional engagement, this study may help explain discrepancies between studies showing obligatory effects of race in very early stages of processing (Ito & Urland, 2005) and those showing flexibility in later stages (Wheeler & Fiske, 2005). We suggest that flexibility in the later effects of race may stem, in part, from goal-based flexibility in the engagement of visual attention.

References

Brosch, T., & Sharma, D. (2005). The role of fear-relevant stimuli in visual search: A comparison of phylogenetic and ontogenetic stimuli. *Emotion*, *5*(3), 360–364.

De Houwer, J. (2003). The extrinsic affective Simon task. *Experimental Psychology (Formerly Zeitschrift Für Experimentelle Psychologie)*, *50*(2), 77–85.

Donders, N. C., Correll, J., & Wittenbrink, B. (2008). Danger stereotypes predict racially biased attentional allocation. *Journal of Experimental Social Psychology*, *44*, 1328–1333.

Fox, E., Russo, R., Bowles, R. J., & Dutton, K. (2001). Do threatening stimuli draw or hold visual attention in sub-clinical anxiety? *Journal of Experimental Psychology: General*, *130*, 681–700.

Ito, T. A., & Cacioppo, J. T. (2000). Electrophysiological evidence of implicit and explicit categorization processes. *Journal of Experimental Social Psychology*, *36*, 660–676.

Ito, T. A., & Urland, G. R. (2003). Race and gender on the brain: Electrocortical measures of attention to race and gender of multiply categorizable individuals. *Journal of Personality & 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*(1), 21–36.

Koster, E. H., Crombez, G., Verschuere, B., & De Houwer, J. (2004). Selective attention to threat in the dot probe paradigm: Differentiating vigilance and difficulty to disengage. *Behaviour research and therapy*, *42*, 1183–1192.

Kurzban, R., Tooby, J., & Cosmides, L. (2001). Can race be erased? Coalitional computation and social categorization. *Proceedings of the National Academy of Sciences*, *98*(26), 15387–15392.

Lang, P. J., Bradley, M. M., & Cuthbert, B. (1997). Motivated attention: Affect, activation and action. In P. J. Lang, R. F. Simons, & M. T. Balaban (Eds.), *Attention and orienting: Sensory and motivational processes* (pp. 97–135). Hillsdale, NJ: Erlbaum, Inc.

Lipp, O. V., & Derakshan, N. (2005). Attentional bias to pictures of fear-relevant animals in a dot probe task. *Emotion*, *5*(3), 365.

Luck, S. J., & Hillyard, S. A. (1994). Electrophysiological correlates of feature analysis during visual search. *Psychophysiology*, *31*, 291–308.

Macrae, C. N., Bodenhausen, G. V., Milne, A.B., Thorn, T. M., & Castelli, L. (1997). On the activation of social stereotypes: The moderating role of processing objectives. *Journal of Experimental Social Psychology*, *33*(5), 471–489.

Mogg, K., Holmes, A., Garner, M., & Bradley, B. P. (2008). Effects of threat cues on attentional shifting, disengagement and response slowing in anxious individuals. *Behaviour Research and Therapy*, *46*(5), 656–667.

Notebaert, L., Crombez, G., Van Damme, S., De Houwer, J., & Theeuwes, J. (2011). Signals of threat do not capture, but prioritize attention: A classical conditioning approach. *Emotion*, *11*, 81–89.

Öhman, A., Flykt, A., & Esteves, F. (2001). Emotion drives attention: Detecting the snake in the grass. *Journal of Experimental Psychology: General*, *130*(3), 466–478.

Öhman, A., Lundqvist, D., & Esteves, F. (2001). The face in the crowd revisited: A threat advantage with schematic stimuli. *Journal of Personality and Social Psychology*, *80*(3), 381–396.

Olsson, A., Ebert, J. P., Banaji, M. R., & Phelps, E. A. (2005). The role of social groups in the persistence of learned fear. *Science*, *309*(5735), 785–787.

Pessoa, L., McKenna, M., Gutierrez, E., & Ungerleider, L. G. (2002). Neural processing of emotional faces requires attention. *Proceedings of the National Academy of Sciences*, *99*(17), 11458–11463.

Pessoa, L., Padmala, S., & Morland, T. (2005). Fate of unattended fearful faces in the amygdala is determined by both attentional resources and cognitive modulation. *NeuroImage*, *28*(1), 249–255.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*, 3–25.

Pourtois, G., Schettino, A., & Vuilleumier, P. (2013). Brain mechanisms for emotional influences on perception and attention: What is magic and what is not. *Biological psychology*, *92*(3), 492–512.

Trawalter, S., Todd, A.R., Baird, A. A., & Richeson, J. A. (2008). Attending to threat: Race-based patterns of selective attention. *Journal of Experimental Social Psychology*, *44*(5), 1322–1327.

Treisman, A. (1964). Selective attention in man. *British Medical Bulletin*, *20*, 12–16.

Treisman, A. (1998). Feature binding, attention and object perception. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *353*(1373), 1295–1306.

Vogt, J., De Houwer, J., & Crombez, G. (2011). Multiple Goal Management Starts With Attention: Goal prioritizing affects the allocation of spatial attention to goal-relevant events. *Experimental Psychology*, *58*(1), 55–61.

Vogt, J., De Houwer, J., Crombez, G., & Van Damme, S. (2013). Competing for attentional priority: Temporary goals versus threats. *Emotion*, *13*(3), 587–598.

Vogt, J., Lozo, L., Koster, E. H. W., & De Houwer, J. (2011). On the role of goal relevance in emotional attention: Disgust evokes early attention to cleanliness. *Cognition and Emotion*, *25*, 466–477.

Vuilleumier, P., & Huang, Y. -M. (2009). Emotional attention: Uncovering the mechanisms of affective biases in perception. *Current Directions in Psychological Science*, *18*, 148–152, http://dx.doi.org/10.1111/j.1467-8721.2009.01626.x.

Wheeler, M. E., & Fiske, S. T. (2005). Controlling racial prejudice social-cognitive goals affect amygdala and stereotype activation. *Psychological Science*, *16*(1), 56–63.

Wolfe, J. M., & Horowitz, T. S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience*, *5*, 1–7.

Young, S. G., Brown, C. M., & Ambady, N. (2012). Priming a natural or human-made environment directs attention to context-congruent threatening stimuli. *Cognition & Emotion*, *26*(5), 927–933.