Journal of Personality and Social Psychology
2004, Vol. 87, No. 6, 876–893

Copyright 2004 by the American Psychological Association
0022-3514/04/$12.00 DOI: 10.1037/0022-3514.87.6.876

Seeing Black: Race, Crime, and Visual Processing

Jennifer L. Eberhardt
Stanford University

Phillip Atiba Goff
The Pennsylvania State University

Valerie J. Purdie
Yale University

Paul G. Davies
University of California, Los Angeles

Using police officers and undergraduates as participants, the authors investigated the influence of stereotypic associations on visual processing in 5 studies. Study 1 demonstrates that Black faces influence participants' ability to spontaneously detect degraded images of crime-relevant objects. Conversely, Studies 2–4 demonstrate that activating abstract concepts (i.e., crime and basketball) induces attentional biases toward Black male faces. Moreover, these processing biases may be related to the degree to which a social group member is physically representative of the social group (Studies 4–5). These studies, taken together, suggest that some associations between social groups and concepts are bidirectional and operate as visual tuning devices—producing shifts in perception and attention of a sort likely to influence decision making and behavior.

The stereotype of Black Americans as violent and criminal has been documented by social psychologists for almost 60 years (Allport & Postman, 1947; Correll, Park, Judd, & Wittenbrink, 2002; Devine, 1989; Duncan, 1976; Greenwald, Oakes, & Hoffman, 2003; Payne, 2001; Sagar & Schofield, 1980). Researchers have highlighted the robustness and frequency of this stereotypic association by demonstrating its effects on numerous outcome variables, including people's memory for who was holding a deadly razor in a subway scene (Allport & Postman, 1947), people's evaluation of ambiguously aggressive behavior (Devine, 1989; Duncan, 1976; Sagar & Schofield, 1980), people's decision to categorize nonweapons as weapons (Payne, 2001), the speed at which people decide to shoot someone holding a weapon (Correll et al., 2002), and the probability that they will shoot at all (Correll et al., 2002; Greenwald et al., 2003). Not only is the association between Blacks and crime strong (i.e., consistent and frequent), it also appears to be automatic (i.e., not subject to intentional control; Payne, 2001; Payne, Lambert, & Jacoby, 2002).

Jennifer L. Eberhardt, Department of Psychology, Stanford University; Phillip Atiba Goff, Department of Psychology, The Pennsylvania State University; Valerie J. Purdie, Department of Psychology, Yale University; Paul G. Davies, Department of Psychology, University of California, Los Angeles.

This research was supported by National Science Foundation Grant BCS-9986128 and a grant from the Research Institute of Comparative Studies in Race and Ethnicity at Stanford University awarded to Jennifer L. Eberhardt. We thank Nalini Ambady, R. Richard Banks, Anders Ericsson, Hazel Markus, Benoit Monin, Jennifer Richeson, Lee Ross, Claude Steele, and Robert Zajonc for their helpful comments on versions of this article.

Correspondence concerning this article should be addressed to Jennifer L. Eberhardt, Department of Psychology, Stanford University, Jordan Hall, Building 420, Stanford, CA 94305. E-mail: jle@psych.stanford.edu

The paradigmatic understanding of the automatic stereotyping process—indeed, the one pursued in all of the research highlighted above—is that the mere presence of a person can lead one to think about the concepts with which that person's social group has become associated. The mere presence of a Black man, for instance, can trigger thoughts that he is violent and criminal. Simply thinking about a Black person renders these concepts more accessible and can lead people to misremember the Black person as the one holding the razor. Merely thinking about Blacks can lead people to evaluate ambiguous behavior as aggressive, to miscategorize harmless objects as weapons, or to shoot quickly, and, at times, inappropriately. In the current article we argue that just as Black faces and Black bodies can trigger thoughts of crime, thinking of crime can trigger thoughts of Black people—that is, some associations between social groups and concepts are bidirectional.

Although contemporary social psychological research has exhaustively documented the fact that social groups can activate concepts (e.g., Bargh, Chen, & Burrows, 1996; Brewer, Dull, & Lui, 1981; Chen & Bargh, 1997; Dovidio, Evans, & Tyler, 1986; Dovidio, Kawakami, Johnson, Johnson, & Howard, 1997; Fazio, Jackson, Dunton, & Williams, 1995; Gaertner & McLaughlin, 1983; Gilbert & Hixon, 1991; Kawakami, Dion, & Dovidio, 1998; Lepore & Brown, 1997; Macrae, Bodenhausen, & Milne, 1995; Macrae, Bodenhausen, Milne, Thorn, & Castelli, 1997; Macrae, Stangor, & Milne, 1994; Moskowitz, Gollwitzer, Wasel, & Schaal, 1999; Perdue & Gurtman, 1990; Wittenbrink, Judd, & Park, 1997), only a small number of studies have probed the converse: the possibility that concepts (by themselves) can activate social groups (Blair & Banaji, 1996; Kawakami & Dovidio, 2001; Kawakami, Dovidio, Moll, Hermsen, & Russin, 2000). In one such study, Blair and Banaji (1996) found that participants exposed to feminine or masculine primes were able to more quickly categorize as female or male those targets consistent with the primes. For instance, after

participants were exposed to such words as *flowers* or *diet,* they categorized female targets faster than male targets. Using the same technique, Kawakami and colleagues (Kawakami & Dovidio, 2001; Kawakami et al., 2000) later demonstrated that Black stereotypic primes could facilitate the racial categorization of Black faces as well. In their studies, stereotypic traits appeared to automatically prime the Black racial category just as the Black racial category automatically primed stereotypic traits.

These results seem perplexing when considered in the context of standard associative network models of stereotyping (Anderson & Klatzky, 1987; Fazio et al., 1995; Lepore & Brown, 1997). The associative network approach suggests that social category nodes will more readily activate concept nodes than the reverse. According to such models, the likelihood that one node will activate the other depends on the strength of the associative link (Fazio, Sanbonmatsu, Powell, & Kardes, 1986; Fazio, Williams, & Powell, 2000; Neely, 1977). Social categories (e.g., Black Americans) tend to be strongly associated with a limited, richly connected set of concepts (e.g., aggressive, musical, athletic, poor). Concepts, in contrast, tend to have broad, sparse associations (Anderson & Klatzky, 1987). For example, the concept "aggressive" is associated with a diverse assortment of social categories, including Black Americans, politicians, panhandlers, stockbrokers, Israelis, athletes, New Yorkers, Italians, men, and so forth. Theoretically, the multiplicity of categories associated with the concept should weaken or dampen the activation of any specific category.

Notwithstanding the large number of social categories that might be associated with a particular concept, bidirectional effects may be especially likely when a specific social category functions as a prototype for a concept. We propose that the Black racial category functions as the prototypical associate for a number of ostensibly race-neutral concepts, such as crime, jazz, basketball, and ghetto. These concepts may trigger clear, visual images of Black Americans. Moreover, not only might the prototypicality of the social category influence the likelihood that the category will be activated by the concept, the activation of the concept may bring to mind prototypical category members. Crime, for example, may trigger images of those Black Americans who seem most physically representative of the Black racial category (i.e., those who look highly stereotypical). Likewise, highly stereotypical Blacks should be the most likely to trigger thoughts of crime.

Explicit consideration of bidirectionality could lead to theoretical refinements of contemporary stereotyping models. Rather than focusing on the capacity of social categories to strongly activate a limited number of concepts, these models might also focus on the capacity of some concepts to strongly activate a limited number of social categories—that is, two routes to maintaining automatic associations could be considered rather than one.[1] Bidirectionality might also help to explain the durability of certain stereotypic associations. Given the existence of two associative routes, automatic associations may be activated and practiced substantially more than previously recognized—even in the absence of initial exposure to a social group member. In a crime-obsessed culture, for example, simply thinking of crime can lead perceivers to conjure up images of Black Americans that "ready" these perceivers to register and selectively attend to Black people who may be present in the actual physical environment.

We argue that visual perception and attention represent core visual practices by which bidirectional associations are reflected and maintained. Bidirectional associations function as visual tuning devices—directing people's eyes, their focus, and their interpretations of the stimuli with which they are confronted. To a large extent, these associations cause people to see (and not to see) in similar ways, despite individual differences in explicit racial attitudes.

We propose that bidirectional associations operate as visual tuning devices by determining the perceptual relevance of stimuli in the physical environment. That is, given the processing capacity limitations that all perceivers face, these associations determine which information is important and worthy of attention and which is not. So, for example, the association of Blacks with crime renders crime objects relevant in the context of Black faces and Black faces relevant in the context of crime. The determination of relevance should have substantial consequences for visual perception and attention in particular. According to our predictions, stimuli deemed relevant should be detected at lower thresholds than stimuli deemed irrelevant. Likewise, attention should be directed toward relevant stimuli and away from irrelevant stimuli.

Of course, the possibility that top-down knowledge influences visual processing has been recognized for quite a long time in the vision sciences (e.g., Goldstein, 1999). Moreover, in contemporary studies, perception researchers are finding evidence for experience-dependent changes in visual processing, even at points in the processing stream that were traditionally thought to be unaffected by top-down information (Dolan et al., 1997; Grill-Spector, Kushnir, Hendler, & Malach, 2000; Kastner, Pinsk, De Weerd, Desimone, & Ungerleider, 1999; Ress, Backus, & Heeger, 2000). Simple manipulations such as instructing participants on where to expect a particular stimulus to appear or allowing participants to practice identifying stimuli at extremely short exposure times can have dramatic effects on visual awareness as well as on neural activation in visual regions of the brain (Grill-Spector et al., 2000; Kastner et al., 1999). The finding that short-term experimental manipulations of this type can tune visual processing may have startling implications for broadly held stereotypic associations between social categories and concepts. Is it possible that these stereotypic associations function as visual tuning devices as well?

Despite recognitions that top-down knowledge modulates a variety of visual processing mechanisms (e.g., shape assignment, figure–ground segregation, object recognition, visual awareness, visual search, attentional selection), empirical demonstrations of *social* influences on vision are rare (e.g., see von Hippel, Sekaquaptewa, & Vargas, 1995). In particular, researchers have not examined how automatic, stereotypic associations can influence object perception when those objects are partially occluded or otherwise degraded. Nor have they examined the influence of such associations on visual attention to faces. Perceiving objects and attending to faces are considered fundamental aspects of vision, and understanding the role of automatic associations could be

[1] The bidirectionality approach we advance here is somewhat reminiscent of the associative symmetry models advanced in cognitive psychology in the late 1950s and 1960s that challenged the assumption of unidirectional effects on cued recall (e.g., see Asch & Ebenholtz, 1962; Horowitz, Norman, & Day, 1966; Horowitz & Prytulak, 1969; Hunt, 1959; Jantz & Underwood, 1958).

critical. Furthermore, we argue that these associations are important not only because they can lead perceivers to make mistakes occasionally but also because they can guide, generally, how perceivers come to organize and structure the visual stimuli to which they are exposed.

Documenting the effects of stereotypic associations on specific visual processing mechanisms could be of great practical significance. For instance, to what extent does seeing Black faces facilitate police officers' detection of guns or knives when they do not have clear images of these objects (e.g., owing to inadequate lighting)? The answer to such a question could significantly improve our understanding of the use-of-force decisions made by police officers. A focus on the bidirectional nature of the Black–crime association places researchers in a position to answer additional questions as well. When ordinary civilians seek to prevent violent crime in their neighborhoods, how likely is it that a Black face will draw their attention? Police officers are routinely faced with the task of solving crime and detecting criminal activity. When police officers are thinking about violent crime, to what extent might they too focus their attention on Black Americans as compared with White Americans? Might Blacks who are most physically representative of the Black racial category be most likely to become the objects of focus? The answers to these questions could have considerable implications for understanding the extent to which both ordinary civilians and police officers engage in racial profiling and why they do so. In fact, these important, practical considerations led us to include both police officers and civilians as study participants in the present research.

Overview of Studies

In the studies that follow, we use a diverse assortment of methods and procedures to more closely examine the association of Blacks and crime and to illustrate its influence on specific visual processing mechanisms. In Study 1, we demonstrate that merely exposing people to Black male faces lowers the perceptual threshold at which they detect degraded images of crime-relevant objects (e.g., guns and knives). In Study 2, we show that exposing people to crime-relevant objects prompts them to visually attend to Black male faces, suggesting that the association of Blacks and criminality is bidirectional. In Study 3, we establish that these effects on visual attention are not simply due to a negative bias toward Blacks; exposing people to a positive concept that has been linked to Blacks leads to similar effects. In Study 4, using different crime primes, different face stimuli, and a slightly different procedure, we demonstrate that activating the crime concept with police officer participants leads them to attend to Black male faces. Moreover, we demonstrate that these crime primes affect officers' memory for the faces to which they were exposed. Priming officers with crime increases the likelihood that they will misremember a Black face as more stereotypically Black than it actually was. Finally, in Study 5, we isolate the association between Blacks and criminality more precisely. When we ask police officers directly, "Who looks criminal?," they choose more Black faces than White faces. The more stereotypically Black a face appears, the more likely officers are to report that the face looks criminal.

Study 1

To demonstrate that bidirectional associations between social groups and concepts influence visual processing, we first sought to establish that exposure to Black faces can decrease the perceptual threshold for recognizing crime-relevant objects. Several recent studies highlight the possibility that the stereotypic association between Blacks and crime influences visual processing (Correll et al., 2002; Greenwald et al., 2003; Payne, 2001). The results from these studies, however, are open to multiple interpretations. For example, Payne (2001) used a sequential priming paradigm to examine the association between Blacks and criminality. He first primed participants with a Black face or a White face on a computer screen and then displayed a gun or a tool. In a forced-choice format, participants were required to indicate with a button push whether the object displayed was a gun or a tool and to do so as quickly as possible. Payne (2001) found that people exposed to Black faces correctly identified guns more quickly than did people exposed to White faces and were more likely, when under time pressure, to misidentify a tool as a gun than were people exposed to White faces. Payne (2001) interpreted this misidentification effect as an automatic perceptual bias. However, in spite of Payne's careful parsing of the misidentification findings, the fact that participants were presented with tools, guns, and faces that were all clearly visible makes it difficult to determine whether participants misidentified harmless objects as guns because they actually "saw" them as guns or because they anticipated seeing guns and so, mistakenly, said that they did. In other words, the race of the face might have produced either a genuine perceptual bias *or* an anticipatory response bias. Indeed, in subsequent research, Payne showed that participants almost always can clearly recognize what the object is, despite the errors they produce (Payne & Shimizu, 2003). Given these results, the extent to which Black faces prompt people to see crime-relevant objects is an issue that warrants further investigation.

In Study 1, we investigated (a) whether the association between Blacks and crime can shift the perceptual threshold for recognizing crime-relevant objects in an impoverished context and (b) whether these perceptual threshold shifts occur despite individual differences in explicit racial attitudes.

To examine this, we subliminally primed participants with Black male faces, with White male faces, or with no faces at all. In a second (ostensibly unrelated) object-detection task, we presented participants with objects on a computer screen that initially were severely degraded and became less degraded in small increments (in 41 picture frames). The participants' task was to indicate (with a button push) the moment at which they could detect what the object was. Importantly, this task did not require a forced choice. Rather, participants simply were asked to write down whatever they thought the object was. The objects were crime relevant (e.g., a gun or a knife) or crime irrelevant (e.g., a camera or a book). Our prediction was that exposure to Black faces would lead to a visual tuning effect, reducing the perceptual threshold for spontaneously recognizing guns and knives, regardless of participants' explicit racial attitudes.

Fazio and colleagues (2000) used a similar procedure to demonstrate the associative strength of object-evaluation associations. Specifically, they primed participants with a category label (e.g., toothpaste) and displayed a brand name (e.g., Colgate) that initially

was degraded but became less so in small increments. The participants' task was to indicate the moment at which they could recognize the brand name. The category label primes facilitated recognition of the brand names. Moreover, association strength predicted the size of the facilitation effect. The greatest facilitation effects emerged for the category label–brand name associations that were most strongly related. Similarly, Macrae and colleagues (1994) demonstrated that social category labels can facilitate the recognition of degraded stereotype-relevant trait words. However, Fazio and colleagues (2000) and Macrae and colleagues (1994) did not use a degraded stimulus procedure to examine how race or crime, in particular, might influence visual processing. In addition, most important, neither examined the extent to which priming might facilitate the detection of real-world objects.

In Study 1, we extend the work of Fazio and colleagues and Macrae and colleagues by examining the extent to which the association between Blacks and crime creates perceptual processing biases that affect object detection. More specifically, Study 1 was intended to directly address the following question: Will activating the Black racial category lower the perceptual threshold for recognizing crime-relevant objects in an impoverished context?

Method

Participants

Participants were 41 White male University of California, Berkeley and Stanford University students who completed the study either for partial course credit or for a $10 payment. To control for potential gender effects, we tested only male students. Participants were contacted via e-mail or through course announcements. Computer error resulted in the loss of data for 2 participants. These participants were excluded from all further analyses, leaving a total of 39 participants.

Design

Participants were randomly assigned to a 3 (race of prime: White prime, Black prime, or no-prime control) × 2 (object type: crime relevant or crime irrelevant) mixed-model design with object type serving as the within-subject factor. The picture frame at which crime-relevant objects could be detected was the primary dependent variable.

Stimulus Materials

Face stimuli. We exposed participants to color photographs of 50 Black or 50 White young adult male faces with neutral facial expressions. The faces were of Stanford students or employees. These photographs were taken from the same face database that we later use for Study 5. The height, weight, age, and attractiveness of the persons photographed did not vary as a function of race. The backgrounds on the photographs were standardized using Adobe Photoshop software.

Object stimuli. We created 14 sets of degraded object stimuli. For each set, a black-and-white line drawing was created of an object. Pixilated "noise" was then added to that image using Adobe Photoshop software, causing the image to look like a television with "snow" or bad reception. Noise was added in equal increments creating 41 picture frames of each object, ranging from an extremely degraded image of the object to a clear image of the object with no degradation added (see Figure 1). These picture frames were then shown in sequence from most degraded (Frame 1) to least (Frame 41). Each frame was presented for 500 ms.

The object stimuli were either crime relevant or crime irrelevant. The crime-relevant objects were line drawings of two guns and two knives. The 10 crime-irrelevant objects were of a pocket watch, a telephone, a bugle horn, a penny, a key, a book, a camera, a cup and saucer, a stapler, and a staple remover. Each crime-irrelevant object was found to be unrelated to crime in pretesting.

Procedure

Participants were scheduled to complete the experiment in pairs. They were greeted by a White experimenter and told that the first task was an "attentional vigilance task." Participants were instructed to focus on a dot at the center of the screen during each trial and were told that "flashes" would appear to the upper and lower left and right of that dot. Participants were seated and the computer monitor arranged such that the flashes appeared 6° from the focus dot. Their goal was to determine (as quickly as possible) whether the flash appeared to the left or the right of the focus dot. The flash consisted of three parts. For participants in the face prime conditions, there was a premask (created from a composite of blurred faces), displayed for 100 ms. Next these participants were exposed to a Black face prime or a White face prime displayed for 30 ms. Last, the postmask (which was identical to the premask) was presented until participants pressed the response key. Participants in the no-prime control condition were presented with the same pre- and postmask, but instead of seeing a face they saw an uninterpretable line drawing produced by Adobe Photoshop software. Participants' detection latency of the flash was mea-

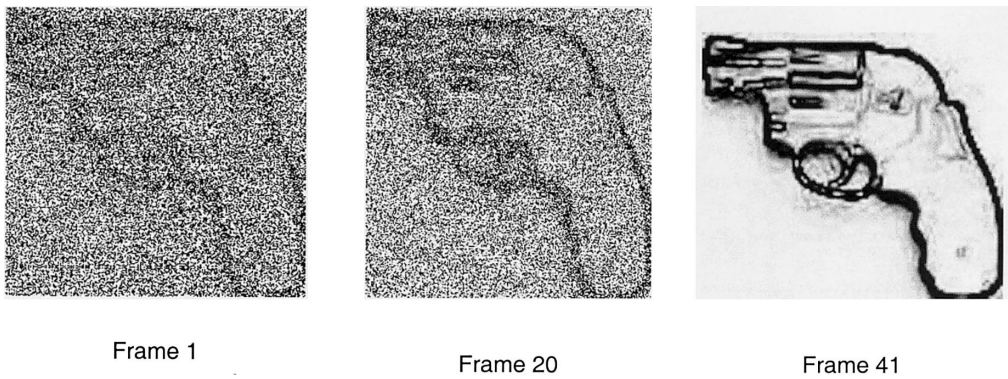

Figure 1. A sample of stimuli used for Study 1. Participants were presented with 41 frames of a continuum displaying an image that initially was severely degraded (e.g., Frame 1), became less degraded (e.g., Frame 20), and finally contained no degradation at all (e.g., Frame 41).

sured from the onset of the postmask to the time participants pressed one of two response keys to indicate that the flash had occurred on either the right or the left side of the screen. Extensive pilot testing revealed that no one was aware of the primes. We exposed participants to the primes subliminally both to reduce suspicion and to reduce the possibility that participants would engage in deliberate strategies to eliminate the effect of the primes on object-detection performance during the second portion of the study. Our priming technique followed closely the paradigm outlined by Bargh and Chartrand (2000).

Participants completed 10 practice trials followed by four blocks of 25 trials, after which the experimenter set up the computer to run the object-detection program. Approximately one third of the participants were subliminally primed with the Black faces during 100% of the "vigilance" task trials, another third were primed with the White faces, and the remaining third were primed with the uninterpretable line drawing.

Participants were told that the second portion of the experimental session would involve an unrelated study on the speed at which people can recognize objects. Participants were told that they would see a series of short "movielike segments" of objects that would start off "fuzzy" and become increasingly easier to identify. Participants were instructed to press the space bar as soon as they knew what the object was. They then had 10 s to write down what the object was. The computer reminded participants when there were 3 s remaining, and participants were thus alerted to the beginning of a new set of presentations. Each participant was exposed to a total of 14 objects (4 crime relevant and 10 crime irrelevant) in this manner. After completing the degraded objects task, participants completed the Modern Racism Scale (MRS; McConahay, 1986) and the Motivation to Control Prejudice Scale (MCP; Dunton & Fazio, 1997), after which they were probed for suspicion, fully debriefed, and thanked for their participation.

Results

Data Reduction

Debriefing responses confirmed that no participants were aware of the primes. Trials in which participants misidentified the object in question were removed. This was a relatively small number of the trials (fewer than 10%). Additionally, there was no effect of race prime on the number or type of errors made ($F < 1$).

Effects of Priming on Object Detection

Of primary interest was the number of picture frames needed to accurately detect the objects as a function of race prime and object type. We expected that participants primed with Black faces would detect crime-relevant images with fewer frames than participants primed with either White faces or no faces. After confirming that the distribution of frames needed to identify an object was not skewed, we submitted the frame data to a 3 (race prime: Black face, White face, or no-prime control) × 2 (object type: crime relevant or crime irrelevant) mixed-model analysis of variance (ANOVA), with object type serving as the within-subject factor. This analysis revealed a significant main effect for race prime, $F(2, 36) = 5.98$, $p < .01$, but no main effect for object type ($F < 1$). As shown in Figure 2, objects presented in the Black face condition ($M = 19.26$) were detected at earlier frames than objects presented in either the no-prime condition ($M = 23.58$) or the White face condition ($M = 24.97$). This main effect, however, was qualified by the predicted Race Prime × Object Type interaction, $F(2, 36) = 7.04$, $p < .01$.

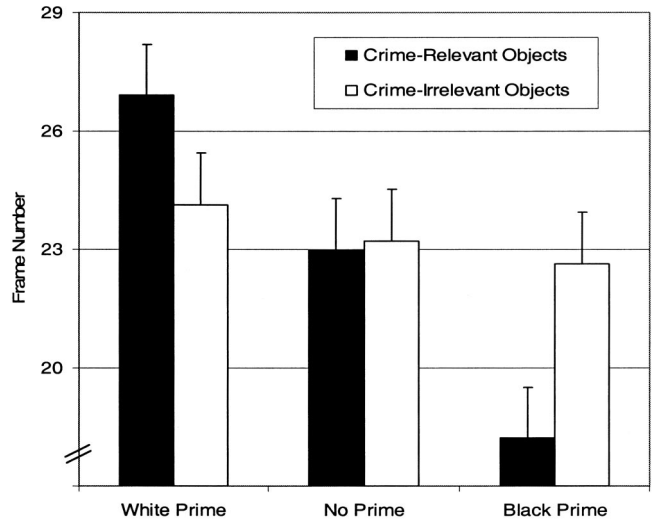

Figure 2. Mean frame number at which the object could be detected as a function of race prime and object type (Study 1). Error bars represent the average standard error for each condition.

As expected, simple effects revealed that in comparison with White face primes, Black face primes dramatically reduced the number of frames needed to accurately detect crime-relevant objects, $t(25) = 4.54$, $p < .01$. Exposure to Black primes also facilitated the detection of crime-relevant objects compared with the no-prime condition, $t(24) = 2.34$, $p < .05$. In contrast, exposure to White primes inhibited the detection of crime-relevant objects compared with the no-prime condition, $t(24) = 2.06$, $p = .05$. As predicted, there was no significant effect of race prime on crime-irrelevant objects ($t < 1$, *ns*).

Participants in the no-prime control condition required the same number of frames to detect crime-relevant and irrelevant objects ($t < 1$, *ns*). After subliminal exposure to Black face primes, however, fewer frames were required to detect crime-relevant objects in comparison with crime-irrelevant objects, $t(13) = 2.96$, $p = .01$. In contrast, after subliminal exposure to White face primes, more frames were required to detect crime-relevant objects in comparison with crime-irrelevant objects $t(12) = 2.35$, $p < .05$.

The Role of Explicit Prejudice

We have argued that stereotypic associations can tune visual perception, regardless of individual differences in explicit prejudice. To measure the potential role of explicit prejudice in producing perceptual threshold shifts, we had participants complete the MRS and MCP after the degraded objects task. Before analyzing participant scores for their potential impact on the frames data, we submitted both MRS scores and MCP scores to a one-way ANOVA to determine whether our priming manipulation had an effect on participants' explicit racial attitudes. This analysis indicated that there was no effect of prime on either MRS or MCP scores (all Fs < 1). We then conducted within-cell correlations between the MRS, the MCP, and our frames. Although some of these correlations were moderate, we found no reliable relationship between participants' explicit racial attitudes and the frame at which they recognized objects (all rs $< .50$, *ns*).

Discussion

The results of Study 1 demonstrate that stereotypic associations have the power to alter the threshold at which real-world objects will be detected. In comparison with White faces, Black faces triggered a form of racialized seeing that facilitated the processing of crime-relevant objects, regardless of individual differences in racial attitudes. Moreover, these results emerged even though participants were not forced to choose between predetermined categories in the degraded objects task.

The results of Study 1 suggest that both Black and White primes tune the detection of crime-relevant objects, yet in opposite directions. Compared with a no-prime control condition, mere exposure to Black faces facilitated the detection of crime-relevant objects. Compared with that same no-prime control condition, however, mere exposure to White faces inhibited the detection of crime-relevant objects. These inhibiting and facilitating effects combined to produce a 21% drop in the perceptual threshold between White and Black face primes (8.8 [White–Black frame difference] / 41 [total number of frames]). Study 1 clearly demonstrates that exposure to racial primes leads to differences in the detection of real-world objects. Studies 2–4 reverse this paradigm by investigating the ways in which real-world objects (i.e., conceptual primes) lead to differences in attention to Black and White male faces.

Study 2

In *The Principles of Psychology,* William James wrote extensively about the power of ideas to direct visual attention. "Attention *creates* no idea," wrote James, "an idea must already be there before we can attend to it. Attention only fixes and retains what the ordinary laws of association bring 'before the footlights' of consciousness" (James, 1890/1950, p. 450). In Study 2, we examine the extent to which Black faces are brought before the footlights of attention when the concept of crime is activated.

Researchers have long shown that schemas influence selective attention (e.g., see Johnson & Dark, 1986). More recently, researchers have shown that the capacity of a stimulus to capture attention depends on the relevance or usefulness of the stimulus to the task that the perceiver is attempting to perform (Yantis & Egeth, 1999). Features of a visual display do not automatically "pop out" in a purely stimulus-driven manner. Rather, attentional capture is, in part, a function of perceiver goals.

Stable personality variables can affect attentional capture as well. For instance, in a classic study by MacLeod, Mathews, and Tata (1986), clinically anxious participants were found to consistently shift their attention toward threat-relevant stimuli whereas nonanxious control participants shifted their attention away from such stimuli. MacLeod and colleagues (1986) introduced a dot-probe paradigm to examine this differential distribution of attention. Pairs of words (threat relevant and neutral) were displayed at different locations on the computer screen. After a brief interval, the words disappeared. A dot appeared immediately afterward, where one of the two words was previously located. The participants' task was to locate the dot as quickly as possible. If participants were more attentive to one of the words, and the dot appeared in its place, they would understandably be relatively quick to find the dot. Dot detection latencies, therefore, were used as a proxy for visual attention. Clinically anxious participants evidenced reduced detection latencies for dot probes in the location of threat-relevant words relative to neutral words. Nonanxious control participants evidenced reduced detection latencies for dot probes in the location of neutral words relative to threat-relevant words. The authors concluded that "high anxiety leads to a bias in selective attention that favors the pickup of emotionally threatening information" (MacLeod et al., 1986, p. 18).

The dot-probe paradigm became the gold standard in personality research on visual attention because it circumvents common problems associated with indirect measures of attentional selection. For example, in Stroop tasks researchers commonly find that participants with high anxiety or depressed affect are slower to name the color of threat-relevant or depression-relevant words (respectively) in comparison with neutral words (see Gotlib, McLachlan, & Katz, 1988; Mathews & MacLeod, 1985). Although this result typically gets interpreted as an attentional bias, it is not clear whether the attentional bias reflects a difference in visual processing. It could, in fact, reflect a difference in how much participants think about threat-relevant or depression-relevant stimuli, independent of visual processing. Unlike the Stroop task, the dot-probe task requires an actual shift in visual attention to a neutral stimulus (dot probe), greatly reducing the possibility that conceptual processing biases or response biases alone could account for the results (see MacLeod et al., 1986).

The dot-probe paradigm has been used now for nearly two decades to directly measure attentional bias of individuals suffering from clinical anxiety and social phobia (Bradley, Mogg, Falla, & Hamilton, 1998; Gilboa-Schechtman, Foa, & Amir, 1999; Maidenberg, Chen, Craske, Bohn, & Bystritsky, 1996; Martin, Williams, & Clark, 1991; Mathews & MacLeod, 1985; Mogg & Bradley, 1999; Mansell, Clark, Ehlers, & Chen, 1999), general dysphoria (Bradley et al., 1997), and depression (Gotlib et al., 1988).

In the current study, we examined the extent to which the association between Blacks and crime would produce an attentional bias toward Black male faces. To measure visual attention, we used a modified version of the dot-probe task used extensively in the personality disorders literature. In what was described as a vigilance task, we activated the concept of crime by subliminally priming participants with crime-relevant objects. Immediately following this priming procedure, participants were introduced to the dot-probe task. During this task, two faces (one Black and the other White) were simultaneously displayed on the computer screen. These faces quickly disappeared and were replaced by a dot probe in the visual location of either face. As in the original MacLeod et al. (1986) study, the participants' task was to locate the dot probe as quickly as possible. We predicted that when the dot probe was placed in the location of the Black face, participants would be faster to detect it when they had been primed with crime than when they had not been primed. The crime prime should render the Black face perceptually relevant, significantly increasing the capacity of this face to capture participants' visual attention. We also examined whether these effects occur beneath awareness. We predicted that participants not only would be unaware that the crime concept had been activated but would also be unaware of where their eyes were attending.

Method

Participants

Participants were 52 White male Stanford University students who completed the study either for partial course credit or for a $7 payment. Participants were contacted via e-mail or through course announcements. Because of a computer failure, data files for 2 participants were incomplete. These participants were excluded from all further analyses, leaving us with a total of 50 participants.

Stimulus Materials

Crime images. Crime-relevant images (used for the priming portion of the study) were chosen on the basis of pretesting conducted in an introductory psychology class. We selected the five most frequently listed images elicited by the question "What images come to mind first when you think of the idea *crime*?", excluding all images that required actual people to be depicted (e.g., mugger or rapist). The images were guns, knives, fingerprints, police badges, and handcuffs. We then created seven line drawings (two of guns, two of knives, and one each of the remaining images) using Adobe Photoshop software. Crime images were displayed one at a time for 30 ms, and each was preceded immediately by a premask and followed immediately by a postmask. We created two dummy images that consisted of jumbled patches from each crime-relevant line drawing. The first dummy image was used as the pre- and postmask. The second dummy image was used as the control prime for participants in the no-prime condition. Extensive pretesting revealed that these dummy images were uninterpretable.

Face stimuli. Ten Black and 10 White faces of clean-shaven men with neutral expressions were pretested for attractiveness. All faces were of individuals with similar height and weight (which we ascertained at the time the photograph was taken via self-report). We then selected 2 Black and 2 White faces that were rated as equally attractive. We digitally standardized the backgrounds in the photographs, leaving just a frontal view of the faces. These faces were then used for the dot-probe portion of the present study.

Vigilance task. We used a priming procedure almost identical to that used in Study 1, with two changes. First, the premask was constructed from jumbled patches of a line drawing, rather than a blurred face. Second, in the prime condition, participants were exposed to crime images, rather than images of Black or White faces.

Dot-probe task. Participants were told that they would participate in a "facial interference" task as the second part of the experiment. They were told that the task intended to measure whether a delay is produced when faces "distract participants" from their task of attentional vigilance. In actuality, this was our dot-probe task, intended to measure attentional bias toward Black or White faces. After two practice trials in which no faces were displayed but, instead, the word *FACE* appeared to the left and right of the focus dot, participants were presented with a focus dot for a randomly determined interval (between 2 and 6 s). One Black and one White face then appeared 6° to the right or left of the dot, with the location of each face randomly determined by the computer. The faces were presented for 450 ms, after which a faint gray dot appeared where one of the two faces was previously located. The location of the dot was, again, randomly determined. Dot-probe latency was measured from the time the target gray dot was displayed to the point at which participants indicated their responses.

Design and Procedure

The study took the form of a 2 (prime: crime prime or no prime) × 2 (dot position: Black face location or White face location) between-subjects factorial design. Dot detection latency served as the primary dependent measure.

Participants were greeted by a White experimenter, and after completing a demographic questionnaire, they participated in two ostensibly separate tasks. Participants were told that the purpose of the first task was simply to gauge how long individuals could remain vigilant to rapidly presented stimuli. This task actually allowed us to subliminally prime participants with crime-relevant images in the primed condition or with a jumbled, uninterpretable image in the no-prime condition. The second task was introduced as a "facial interference" task. Participants were told that the purpose of this task was to examine whether seeing faces would affect one's ability to respond quickly and accurately to stimuli. This second task was actually the dot-probe task.

After participants completed the dot-probe task, they were asked to complete an experimental packet. The packet contained questions about how they were feeling about the study and how they were feeling generally, as well as questions about where the participants thought they were looking. Specifically, participants answered the following two questions regarding attentional awareness: "Which face did you look at first?" and "Which face did you look at longest?" For both questions, participants indicated their answers by circling either "the face on the left" or "the face on the right." Participants were then probed for suspicion, fully debriefed, and thanked for their participation.

Results

Data Transformation

We submitted our detection latency data to a reciprocal transformation, which allowed us to eliminate the positive skew of the data.[2] All subsequent analyses were performed on the transformed data. Because the pattern of means was nearly identical, however, we present the raw detection latencies in Figure 3 for ease of interpretation.

Effects of Priming on Visual Attention

We submitted the transformed detection latencies to a 2 (prime: crime prime or no prime) × 2 (dot position: Black face location or White face location) between-subjects ANOVA. As predicted, this analysis revealed a significant interaction of Prime × Dot Position on detection latency, $F(1, 46) = 11.89$, $p < .01$.

Analysis of the simple effects revealed that the pattern of this significant interaction matched our hypotheses. We predicted that when the dot probe was in the Black face location, participants primed with the crime-relevant images would be faster to find the dot than participants who were not primed. A simple effects test confirmed this specific hypothesis, $F(1, 46) = 8.22$, $p < .01$. The activation of the crime concept, indeed, facilitated the visual pickup of Black male faces. In contrast, we predicted that when the dot was in the White face location, the crime prime would not facilitate dot detection. A simple effects test confirmed this hypothesis as well. When the dot was in the White face location, participants primed with crime were no faster to detect the dot probe than those not primed with crime. In fact, in this situation, participants primed with crime were significantly slower to detect the dot probe than those not primed, $F(1, 46) = 3.91$, $p = .05$. Finally, although participants were significantly faster to find the

[2] A natural log transformation and a square root transformation failed to eliminate the positive skew of the data. Bargh and Chartrand (2000) have recommended that when this occurs, a reciprocal transformation should be used to enable a valid analysis of the data.

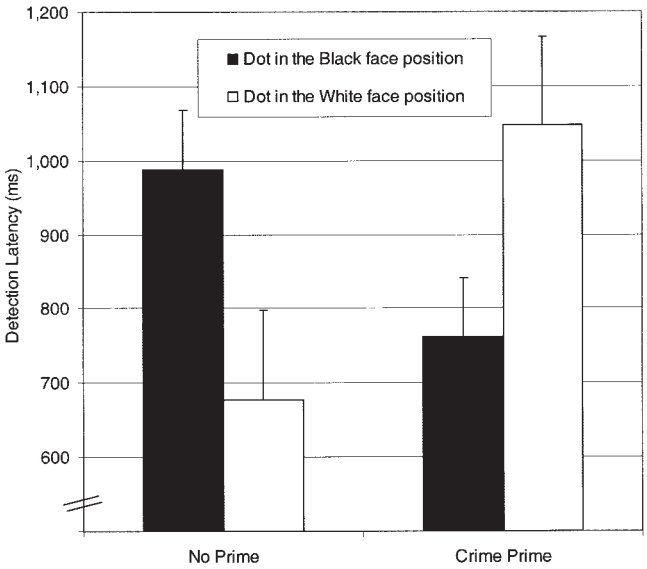

Figure 3. Mean detection latency as a function of prime and dot location (Study 2). Error bars represent the average standard error for each condition.

dot probe in the White face location than in the Black face location when there was no crime prime, $F(1, 46) = 12.02$, $p < .01$, this attentional difference disappeared when participants were primed with crime, $F(1, 46) = 2.07$, $p = .15$. In fact, the pattern reversed.

Participant Awareness of Attentional Biases

Debriefing responses confirmed that no participants were aware of having seen the crime-relevant images. We were most interested, however, in participants' awareness of where they were looking during the dot-probe task. To determine whether participants were aware of their attentional biases, we examined the extent to which participants' reports of where they were looking correlated with detection latencies. The logic behind this is as follows. If (for instance) participants accurately reported looking to the left of the screen, then they would be faster to find the dot when it appeared to the left of the screen. However, if participants were not able to accurately report their attentional bias, then when the dot appeared to the left of the screen, participants who reported looking to the left of the screen would be no faster to find the dot than participants who reported looking to the right. A series of *t* tests revealed no significant relationship between where participants thought they were looking (first or longest) and their attentional bias (as measured by detection latency) (all ts < 2, *ns*).

Subsequently, we performed a *t* test to determine whether participants were aware of an attentional bias toward the Black face (a similar logic governed this analysis). These analyses indicated that participants were not aware of any attentional biases as a result of the prime (all ts < 1). Taken together, this provides evidence that attentional biases produced by stereotypic associations can be unintentional and manipulated beneath awareness.

Discussion

This study provides additional support for the visual tuning hypothesis using a completely different paradigm. The concept of crime affected selective attention such that participants were over 350 ms faster to direct their attention to the location of the Black male face when the concept of crime was activated than when it was not. The results of this study are also consistent with our bidirectionality claim. These results reveal that ostensibly race-neutral concepts such as crime can become racialized. Not only are Blacks thought of as criminal, but also crime is thought of as Black.

The increased visual attention to Black faces brought about by the crime prime is somewhat reminiscent of the phenomenon of "high visibility" that the novelist Ralph Ellison highlighted in his 1950s American classic, *Invisible Man* (Ellison, 1952). Ellison described the Black American predicament as one where Blacks are visually registered only with the aid of cultural stereotypes that function to distort their image. Cultural stereotypes lead Blacks to be the subject of gaze, yet at the same time, these same stereotypes prevent Blacks from being fully seen. Ironically, then, high visibility is accompanied by invisibility. In an Ellisonian sense, here we have shown that Black faces were much more likely to capture the attentional systems of those who had been induced to think about crime than those who had not. It is as if the stereotypic association between Blacks and crime rendered these faces more perceptually relevant and therefore worthy of gaze.

Study 3

We have argued that specific stereotypic associations influence visual processing mechanisms. The bidirectional association between Black Americans and criminality, in particular, can produce perceptual threshold shifts and direct how attention is deployed. Nonetheless, the results of Studies 1 and 2 also could be due to a simple out-group negativity effect. That is, out-groups (such as Blacks) may become associated with any negative concept, regardless of the concept's relevance to the specific stereotypes associated with those groups. Indeed, there are numerous studies demonstrating that Blacks are associated with negatively valenced words and concepts of all types (Dovidio et al., 1997; Fazio et al., 1995; Gaertner & McLaughlin, 1983; Greenwald, McGhee, & Schwartz, 1998). Researchers typically find that Black primes (faces, names, or labels) lead to faster responses to negative words and concepts than White primes.

Study 3 was designed to examine the extent to which the socially induced attentional biases observed in Study 2 are due to concept valence as opposed to concept content. To examine this, we primed participants with a positive concept associated with Black Americans (or not) and measured the effects on attentional deployment. We predicted that the activation of a concept associated with Blacks would lead to an attentional bias for Black American faces, even when the concept activated was positive. We also included measures of explicit prejudice. Once again, we predicted differences in attentional deployment despite individual differences in explicit racial attitudes.

Method

Participants

Participants were 75 White male Stanford University students who completed the study either for partial course credit or for a $7 payment. Participants were contacted via e-mail or through course announcements.

Because of computer malfunctioning, data for 4 participants were incomplete. Additionally, 2 participants evidenced some knowledge of the study's hypothesis. These participants were excluded from all further analyses, leaving a total of 69 participants.

Materials

For the current study, we primed participants with the concept of basketball (or not). The results of a pilot study confirmed that, as with crime, everyone (30 of 30 participants in our pilot sample) has knowledge of an association of Blacks and athletics. Unlike crime, however, the athletic stereotype is positively valenced. In fact, in our pilot study we found the athlete stereotype to be more positive than any other stereotype of Blacks.

In an effort to broaden stimulus sampling, in Study 3 we chose to prime participants with words rather than images. Pilot data clearly demonstrated that of all sports, Blacks are most highly associated with basketball. Specific words relevant to basketball were chosen on the basis of pretesting conducted in an introductory psychology class. For the current study, we selected the 24 most frequently listed words elicited by the question "What words come to mind first when you think of the idea *basketball*?" The words were *assist*, *backboard*, *bankshot*, *basket*, *dribble*, *dunk*, *fastbreak*, *fingerroll*, *freethrow*, *frontcourt*, *fullcourt*, *halfcourt*, *hookshot*, *hoop*, *jumper*, *layup*, *NBA*, *rebound*, *rim*, *shotclock*, *slam*, *swish*, *tip off*, and *traveling*. Basketball-relevant words were displayed one at a time for 75 ms, and each was preceded immediately by a premask and followed immediately by a postmask. Because we did not use images as our subliminal primes, the pre- and postmasks were a nonsense letter string. Participants in the no-prime condition saw a second letter string rather than basketball-relevant words.

Procedure and Design

The study took the form of a 2 (prime: basketball prime or no prime) × 2 (dot position: Black face location or White face location) between-subjects factorial design. Dot detection latency served as the primary dependent measure. Study 3 followed the exact protocol of Study 2 with the following three exceptions: (a) When participants were primed, they were primed with basketball-relevant words instead of crime-relevant images; (b) participants completed the MRS and MCP after completing the dot-probe task; and (c) at the conclusion of the study, participants were not probed on where they looked during the dot-probe task.

Results

Data Transformation

As in Study 2, we first submitted our detection latency data to a reciprocal transformation, which allowed us to eliminate the positive skew of the data. All subsequent analyses were performed on the transformed data. However, because the pattern of means was nearly identical, we present the raw detection latencies in Figure 4 for ease of interpretation.

Effects of Priming on Visual Attention

Debriefing results confirmed that no participants were aware of the basketball primes. We submitted the transformed detection latencies to a 2 (prime: basketball prime or no prime) × 2 (dot position: Black face location or White face location) between-subjects ANOVA. As anticipated, this analysis revealed a significant interaction of Prime × Dot Position on detection latency, $F(1, 65) = 5.33$, $p < .05$. Recall that our primary hypothesis was that when the dot probe was in the location of the Black face, participants primed with the basketball-relevant words would be faster to find the dot than participants who were not primed. A simple effects test confirmed this specific hypothesis, $F(1, 65) = 4.96$, $p < .05$.

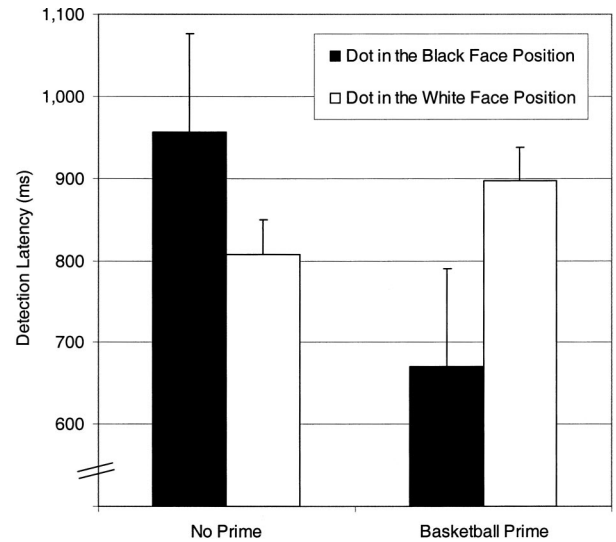

Figure 4. Mean detection latency as a function of prime and dot location (Study 3). Error bars represent the average standard error for each condition.

Although our primary hypothesis was confirmed, we obtained additional findings that were not entirely consistent with the findings from Study 2. Though the pattern of means was identical to that of Study 2, participants who saw the dot appear in the White face location were not significantly slower to find it when primed with basketball-relevant words than when they were not primed ($F < 1$). Also, the attentional bias toward the White face in the unprimed condition was not significant ($F < 1$). Priming participants with basketball-relevant words, however, did produce a significant attentional bias toward the Black face. Though participants showed no significant attentional bias toward either face when they were not primed, they were significantly faster to find the dot in the Black face location than in the White face location when primed with basketball-relevant words, $F(1, 65) = 6.60$, $p = .01$.

The Role of Explicit Prejudice

To measure the role of explicit prejudice in producing attentional bias, we had participants complete the MRS and MCP after the dot-probe task. Before analyzing participant scores for their potential impact on detection latency data, we submitted both MRS scores and MCP scores to the same 2 × 2 ANOVA as our detection latency data. This was done to ensure that participants' prejudice scores were not influenced by our manipulations. The two univariate ANOVAs yielded no reliable main effects or interactions, all $Fs(1, 65) < 3.00$, *ns*. We then conducted within-cell correlations between the MRS, the MCP, and our detection latencies. These correlations revealed no reliable relationship between

participants' explicit racial attitudes and their reaction times (all rs $< .40$, *ns*).

Discussion

The results of Study 3 demonstrate that stereotypic associations other than crime can lead to visual tuning effects. When the dot probe was in the location of the Black face, basketball-primed participants located it faster than unprimed participants. This result confirmed our primary hypothesis and was identical to the result we obtained in Study 2 with the crime prime. Although the overall pattern of results is quite similar to the results of Study 2, there were discrepancies with some of the additional findings that require discussion. For example, the bias to attend to White faces rather than Black faces in the no-prime condition was not significant in Study 3. Therefore, in comparison with Study 2, the baseline result was shifted. In Study 2, the prime wiped away the significant White face bias present in the no-prime, baseline condition. In Study 3, the prime significantly reversed the nonsignificant White face bias present in the baseline condition. The pattern of results is the same across Studies 2 and 3. However, the results in the initial baseline condition are stronger in Study 2 than in Study 3, and this may have produced the difference in the magnitude of the Black face bias in the prime condition across the studies. Alternatively, it is possible that the larger Black face bias in the prime condition in Study 3 occurred because we primed participants in Study 3 with precise words relevant to the concept rather than images that arguably are more vague and open to multiple interpretations.

Nevertheless, Study 3 extends Study 2 by demonstrating that stereotypes can influence visual attention regardless of their valence. Crime and basketball produced changes in attentional deployment because both concepts are strongly associated with Black Americans. In this study, as in Study 2, stereotypic associations pushed Black faces into the "footlights" of attention. Moreover, this study demonstrates that these strong associations produced changes in attentional deployment, regardless of individual differences in explicit racial attitudes. We offer the results of Study 3 as additional evidence that the more specific link between Blacks and criminality led to the visual tuning effects observed in Studies 1 and 2, rather than a general negative bias directed at Blacks.

Study 4

We now know that both positive and negative stereotypic associations can tune visual attention. However, we know less about how such visual tuning effects are accomplished. For example, do stereotypic associations affect attentional deployment by determining where participants look first or by determining where they look longest? Both dot-probe studies used thus far displayed the Black and White faces for 450 ms, and the onset of the dot-probe immediately followed. Is it the case that the stereotypic association between Blacks and crime led participants to look at the Black face in the first half-second only, or did this association also cause their attention to linger on the Black face?[3] Moreover, how might attentional bias affect people's memory for the faces displayed? Do stereotypic associations lead to a more accurate visual memory, or might they lead to stereotype-consistent distortions in memory? As attentional bias works to "fix and retain" a visual image (to use William James's words), what precisely is the image that is fixed and retained? Will activating the crime concept cause participants to retain an image of a Black face that is highly stereotypically Black and thus strongly representative of the Black racial category? Study 4 was designed to begin to answer some of these questions regarding the mechanisms of attentional bias.

Study 4 was also designed to begin examining the extent to which stereotype-induced attentional biases generalize to other participant populations. For example, might such visual tuning effects extend to police officers, who are charged with the task of investigating criminal activity? Whose faces do they seize upon when they think of capturing, shooting, arresting, or apprehending? We strongly suspect that the pattern of attentional deployment we have obtained with undergraduates will generalize to police officers, because police officers have knowledge of the very same stereotypic associations as the undergraduates we have tested. Alternatively, one might imagine that police officers would not exhibit an attentional bias in the dot-probe paradigm, given their high level of training and sophisticated knowledge about crime. Perhaps even still, one might imagine that they would exhibit an exaggerated form of attentional bias, given their experiences with Blacks in the specific context of crime in addition to their exposure to the general stereotypic association of Blacks and crime.

Study 4 extends Studies 2 and 3 along several critical dimensions. To examine the extent to which visual tuning effects generalize to other participant populations, we used police officers as study participants. To examine the extent to which a stereotypic association can both direct attention to a specific location and allow attention to linger in that location, we systematically varied the duration at which the Black and White faces were displayed (from 450 ms to 650 ms to 850 ms); and finally, to examine the potential effects of attentional bias on visual memory, we gave officers a surprise face-recognition task just before the conclusion of the study.

Method

Participants

Sixty-one police officers from a police department voluntarily participated in this study. The police department, which is located in an urban setting in the United States, provides services to well over 100,000 civilians. In this agency, 76% of the officers are White, 86% are male, and the average age is 42. We did not collect data on police officer gender or race for this study. Data from 2 officers were excluded because of a computer malfunction. One officer's data were excluded because the officer did not complete the study, and one officer's data were excluded because of the large number of errors the officer committed on the vigilance task (more than 2 standard deviations higher than the mean of the sample). This left us with 57 participants of the original 61, on which all analyses were conducted.

Materials

Crime primes. Crime words were chosen as primes (rather than crime images) in an effort to broaden stimulus sampling and to tailor the primes more precisely to the specific participant population used in the study. To this end, police officer participants were primed with words associated with

[3] We thank Russell Fazio and Marilyn Brewer for raising this point.

enforcing the law against violent criminals. After polling police officers, we chose the following 10 words to serve as primes (because they were the words most commonly listed): *violent*, *crime*, *stop*, *investigate*, *arrest*, *report*, *shoot*, *capture*, *chase*, and *apprehend*.

Face stimuli. In an additional effort to broaden stimulus sampling, target faces were chosen from a database of prisoners who were convicted of first-degree murder in the state of Florida. These prisoners' faces were pretested on attractiveness and stereotypicality. Pilot participants were not told that these faces were the faces of convicted criminals. Pilot participants who were instructed to rate the stereotypicality of the faces were told that they could use any number of physical features (e.g., the lips, the nose, the hair texture, the skin tone) to make such a judgment. They were asked to look at a series of 60 Black male faces and to use the physical features that most people commonly associate with Blacks to provide us with a stereotypicality rating of each face. A second group of pilot participants were shown a series of 60 White male faces and were asked to use the physical features people commonly associate with Whites to provide us with a stereotypicality rating.

After receiving these ratings, we chose five faces within each race, one from each quintile of the stereotypicality distribution provided by the pilot participants. Each face was also matched for attractiveness across race. A Black face lineup and a White face lineup were then created that included a target face along with four additional faces. Within each face lineup, two faces were less stereotypical than the target (i.e., from the first and second quintiles) and two faces were more stereotypical than the target (i.e., from the fourth and fifth quintiles). The Black and White target faces were selected from the middle quintile of the stereotypicality distribution. These Black and White lineups were later used during the surprise face-recognition task.

Procedure and Design

The study took the form of a 2 (prime: crime prime or no prime) × 2 (dot position: Black face location or White face location) between-subjects factorial design. Police officers were tested on site at the police department in small groups ranging from 2 to 5 participants. Study 4 followed the exact protocol of Study 2 with the exception of the changes to the crime primes, the face stimuli, the presentation duration of the face stimuli in the dot-probe task (this varied from 450 ms to 650 ms to 850 ms across participants), and the inclusion of a surprise face-recognition memory task.

Participants were given the surprise face-recognition memory task after they completed the dot-probe task. Participants were exposed to a Black face lineup and a White face lineup. For each lineup, participants were asked to identify the face that had been displayed during the dot-probe task. For each lineup, all five faces of one race—the target and four distracters—were presented on the computer screen simultaneously. The order in which participants saw the Black and White lineups was randomly determined, as was the location of each face on the screen. Participants were asked to indicate their choice in the first lineup, then the second lineup, and were then debriefed.

We were interested in the degree to which the crime prime would influence officers' memories for the original target faces displayed during the dot-probe task. To the extent that the face-recognition memories of the officers were inaccurate, we were poised to examine whether the crime prime influenced the pattern of errors. We were especially interested in whether officers would be more likely to falsely identify a face from the Black lineup that was more stereotypical than the actual target face when they were primed with crime than when they were not.

Results

Data Transformation

As in Studies 2 and 3, we first submitted our detection latency data to a reciprocal transformation, which allowed us to eliminate the positive skew of the data. All subsequent analyses were performed on the transformed data. However, because the pattern of means was nearly identical, we present the raw detection latencies in Figure 5 for ease of interpretation.

Effects of Face Presentation Duration

A one-way ANOVA indicated that dot detection latencies were not significantly influenced by the amount of time each face was displayed ($F < 1$). Similarly, including face presentation duration as a covariate in subsequent analyses of dot detection latency and face stereotypicality yielded no significant results (Fs < 1), nor did it impact the analyses in which it was included. Consequently, we collapsed across face presentation duration in the remaining analyses.

Effects of Priming on Visual Attention

Debriefing results confirmed that no participants were aware of the primes. We submitted the transformed detection latencies to a 2 (prime: crime prime or no prime) × 2 (dot position: Black face location or White face location) ANOVA. As anticipated, this analysis revealed a significant interaction of Prime × Dot Position on detection latency, $F(1, 53) = 15.24$, $p < .01$. Recall that our primary hypothesis was that when the dot probe was in the location of the Black face, officers primed with the crime-relevant words would be faster to find the dot than officers who were not primed. A simple effects test confirmed this specific hypothesis, $F(1, 53) = 3.95$, $p < .05$. In addition, a simple effects test confirmed that officers primed with crime were slower to find the dot behind the White face than officers who had not been primed with crime, $F(1, 53) = 12.60$, $p < .01$.

There was also an attentional bias toward the White face in the unprimed condition such that participants found the dot faster when it was in the White face location than the Black face location, $F(1, 53) = 9.74$, $p < .01$. Moreover, priming participants with

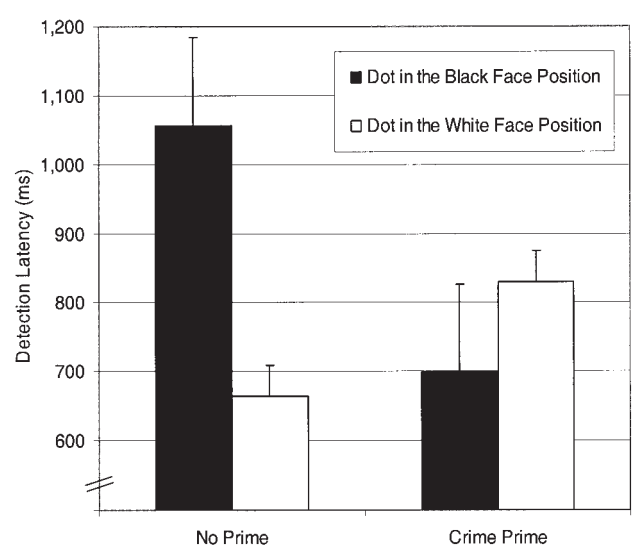

Figure 5. Mean detection latency as a function of prime and dot location (Study 4). Error bars represent the average standard error for each condition.

crime-relevant words reversed this relationship, such that participants found the dot faster when it was in the Black face location than when it was in the White face location, $F(1, 53) = 5.87$, $p < .05$.

Error Rates During the Memory Task

There was no overall difference in error rates on the face-recognition memory task as a function of the prime ($F < 1$). The average accuracy rate was 34%, which was significantly above chance (1 in 5), $t(56) = 10.49$, $p < .01$.

Stereotypicality Ratings of Faces Identified in the Memory Task

Each face in the Black and White lineups was coded in terms of the stereotypicality quintile from which it was taken. The faces taken from the lowest quintile were coded as –2, the next least stereotypical faces were coded as –1, the targets were coded as 0, the faces in the next quintile were coded as 1, and the most stereotypical faces were coded as 2. The data were then subjected to a 2 (race of face: Black or White) × 2 (prime: crime prime or no prime) mixed-model ANOVA with race of face as the within-subject variable. We were primarily interested in the extent to which the crime prime would produce false identifications in the Black lineup such that faces more stereotypically Black than the target would be mistaken for the target.

Our analysis revealed a reliable main effect of race of face, such that participants identified more stereotypically Black faces ($M = .46$) than stereotypically White faces ($M = -.30$), $F(1, 55) = 16.82$, $p < .01$. One-sample t tests further revealed that participants reliably identified faces that were more stereotypically Black than the Black target, $t(56) = 3.03$, $p < .01$. There was a marginally significant trend in the opposite direction for White faces, $t(56) = 1.76$, $p = .08$. This main effect, however, was qualified by a two-way interaction, $F(1, 55) = 7.30$, $p < .01$. Simple effects tests revealed that participants indeed chose more stereotypically Black faces as targets when primed with crime ($M = .81$) than when not primed ($M = .13$), $t(55) = 2.35$, $p < .05$. There was not, however, a significant effect of stereotypicality for White faces as a function of the prime ($t = 1$, *ns*). These means are shown in Figure 6.

Discussion

As predicted, police officers exhibited the same pattern of attentional bias as the undergraduate participants in Studies 2 and 3. These results with police officers were obtained despite changes in the crime primes, the face stimuli, and the face presentation duration. For example, not only did the crime primes influence where officers looked at the initial point of measurement (i.e., at 450 ms), these primes also influenced where officers continued to look. Moreover, when officers were tested on their memory for the target faces, they were more likely to falsely identify a face that was more stereotypically Black than the target when they were primed with crime than when they were not primed. It appears as though stereotypic associations led perceivers to look in a partic-

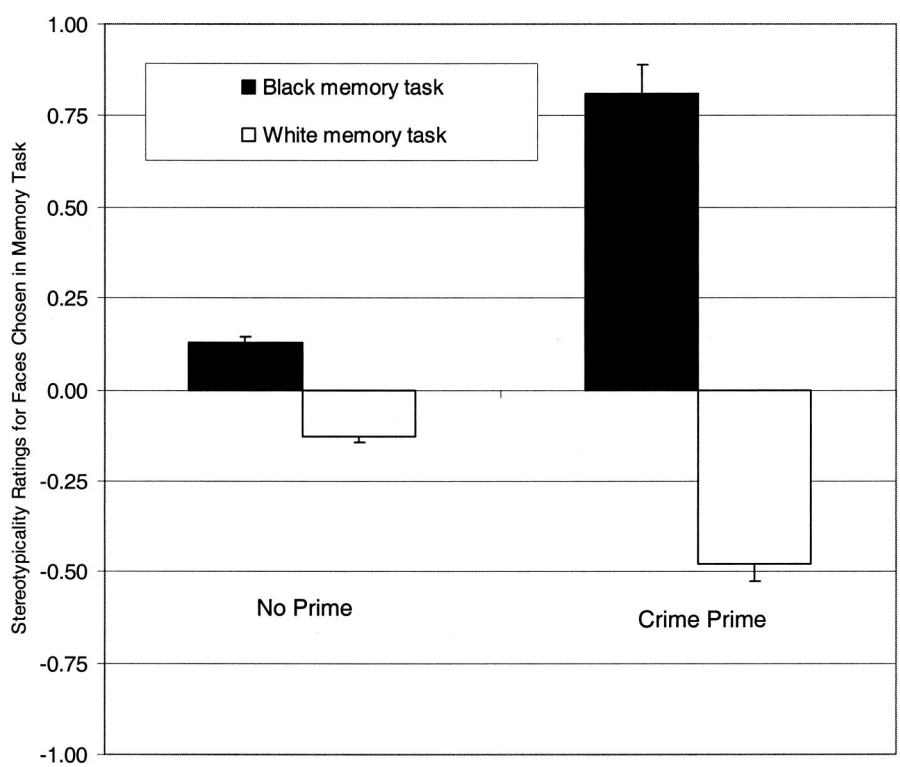

Figure 6. Mean stereotypicality of faces identified in memory task as a function of prime and race (Study 4). Error bars represent the average standard error for each cell.

ular location, yet what perceivers were able to remember was, in part, a function of these stereotypic associations. That is, priming police officers with crime caused them to remember Black faces in a manner that more strongly supports the association between Blacks and criminality. When these officers were asked, "Which face did you see?," priming them with crime led them to envision a Black face that was even more strongly representative of the Black racial category than the Black face to which they were actually exposed. Thus, thoughts of violent crime led to a systematic distortion of the Black image—a phenomenon that Ralph Ellison so masterfully highlighted over 50 years ago.

Generally, these results are significant because they suggest that the process of visually attending to a stimulus will not always aid perceptual memory. These results also are significant, however, because they demonstrate the influence of strong, stereotypic associations on face processing mechanisms in particular (see also Eberhardt, Dasgupta, & Banaszynski, 2003). Practically, this could have implications for eyewitness testimony. For example, Blacks who appear most stereotypically Black may be most vulnerable to false identifications in real criminal lineups. This type of false identification may be likely even when the actual perpetrator is present in the lineup and even when the eyewitness was visually drawn to the perpetrator's face at the time of the crime.

These results also may provide a unique demonstration of association strength. We have argued that association strength increases not only the likelihood that social categories will trigger concepts but also the likelihood that a concept will trigger a social category. We now have some initial evidence that exposure to a concept can lead to the triggering of a social category image that is strongly representative of the social category. Indeed, thinking about the concept of crime not only brought Black faces to mind but brought stereotypically Black faces to mind.

Study 5

Study 5 examines directly an assumption on which our discussion of the memory results from Study 4 was premised: Police officers view more stereotypically Black faces as more criminal. To examine this, in Study 5 we presented police officers with Black and White male faces and asked the question, "Who looks criminal?" We predicted that police officers would choose more Black faces than White faces as criminal and that Black faces rated high in stereotypicality would be even more likely to be perceived as criminal than Black faces rated low in stereotypicality. In other words, we predicted that police officers would use the physical features linked to race to inform them about who looks criminal. Recently, researchers have documented that people are attentive to physical trait variation among Black Americans (Blair, Judd, Sadler, & Jenkins, 2002; Livingston, 2001; Maddox & Gray, 2001; Williams & Eberhardt, 2004). Here we argue that police officers imbue this physical variation with criminal meaning—that is, the "more Black" an individual appears, the more criminal that individual is seen to be.

Method

Participants

One hundred eighty-two police officers (159 male, 23 female) voluntarily participated in this study. The officers were drawn from the same police department used in Study 4. The racial composition of our sample was as follows: 115 White Americans, 8 Black Americans, 6 Asian Americans, and 1 Native American (52 officers did not disclose their race). Sixteen officers were excluded from the final analysis because they did not follow instructions, leaving a total sample size of 166 officers.

Stimulus Materials

Participants were exposed to color photographs of 40 Black or 40 White male faces (with neutral facial expressions) ranging in age from 18 to 40 years. In this study, the photographs were of male students and employees of Stanford University. The backgrounds on the photographs were standardized using Adobe Photoshop software. These stimuli were then converted to slides and projected onto a screen using a standard slide projector. The projected images were approximately 51 × 41 cm (such that they could be seen clearly by each study participant).

Procedure

The study was conducted on site at the police department in a large room equipped with tables, chairs, and a large screen for stimulus viewing. Officers were tested in small groups of 5–15 participants. On arrival, officers were greeted by two White experimenters, who led them to sit at designated tables. After a lieutenant introduced the experimenters, the experimenters informed the officers that they were conducting a study on face perception. The experimenters then informed the officers about their confidentiality and rights to refuse to participate in the study. All officers gave verbal consent to participate.

Officers were asked to view a series of faces (all of the same race and age group) and to make judgments about them. Approximately half of the participant groups were shown a series of Black male faces, and the remaining half were shown a series of White male faces. The faces appeared on a screen at the front of the room one at a time. Each face appeared for approximately 5 s. Approximately one third of the officers in each participant group completed a stereotypicality measure. These participants were asked to indicate on a 7-point scale (1 = *not at all stereotypical*, 7 = *extremely stereotypical*) how stereotypically Black or White each face stimulus appeared as it was projected onto the screen. These participants were instructed to look at the faces and to use the physical features that most people commonly associate with Blacks (or Whites) to provide us with a rating. Another third of the officers completed a criminality measure. These officers were informed that some of the faces they were about to see might be of criminals. For each face presented, their task was to indicate (by circling yes or no) whether they thought the person "looked criminal." The remaining third of the officers completed an attractiveness measure. These participants were asked to indicate on a 7-point scale (1 = *not at all attractive,* 7 = *extremely attractive*) the extent to which others would find each face attractive as each face stimulus was projected onto the screen. The attractiveness measure was of no theoretical interest but rather allowed us to ensure that the Black and White faces were equated on perceived attractiveness. Each officer completed one measure only. The particular measure an officer completed was randomly determined, with the restriction that each of the three measures would be completed by one third of the officers in any one small group of participants. After completion of the measures, the officers in the group were debriefed, thanked for their participation, and dismissed.

Results

After confirming that the Black (M = 3.45) and White (M = 3.43) faces were perceived as equal in attractiveness ($F < 1$), we dropped attractiveness in all subsequent analyses and turned to our primary interest: How might race and stereotypicality affect judgments of criminality? To examine this, we used faces as our unit

of analysis. Specifically, we averaged officers' ratings of individual faces such that each face had a stereotypicality and criminality rating. Two faces (representing 2.5% of the data) were designated as outliers on stereotypicality (over 2 standard deviations above the mean) and removed from further analysis. Next, we conducted a median split on the stereotypicality data across Black and White faces, yielding two groups: high and low stereotypicality.[4] We then submitted the criminality data to a 2 (race: Black or White) × 2 (stereotypicality: high or low) between-faces ANOVA. This analysis revealed no main effect for stereotypicality on judgments of criminality ($F < 1$). However, as shown in Figure 7, a significant main effect for race emerged, $F(1, 76) = 6.35$, $p = .01$. As predicted, more Black faces ($M = 11.95$) were thought to look criminal than White faces ($M = 9.65$). This race main effect was qualified by a significant Race × Stereotypicality interaction, $F(1, 74) = 4.60$, $p < .05$. As predicted, analysis of simple effects revealed that more Black faces rated high in stereotypicality were judged as criminal ($M = 12.95$) than Black faces rated low in stereotypicality ($M = 10.83$), $F(1, 36) = 4.78$, $p < .05$. This pattern did not emerge for White faces rated high in stereotypicality ($M = 8.80$) in comparison with White faces rated low in stereotypicality ($M = 10.5$), $F(1, 38) = 1.34$, *ns*. Additionally, significantly more Black faces rated high in stereotypicality were judged as criminal ($M = 12.95$) than White faces rated high in stereotypicality ($M = 8.80$), $F(1, 38) = 9.74$, $p < .01$. Finally, a planned contrast analysis revealed that highly stereotypical Black faces were more likely to be judged criminal than any other group in the study, $F(1, 74) = 8.12$, $p < .01$.

Discussion

When officers were given no information other than a face and when they were explicitly directed to make judgments of criminality, race played a significant role in how those judgments were made. Black faces looked more criminal to police officers; the more Black, the more criminal. These results provide additional evidence that police officers associate Blacks with the specific concept of crime. Moreover, these results shed light on the face-recognition memory errors made by police officers in Study 4. In that study, police officers were more likely to falsely identify a Black face that was more stereotypically Black than the target when primed with crime than when not primed with crime. Thinking of crime may have led officers to falsely identify the more stereotypically Black face because more stereotypically Black faces are more strongly associated with the concept of crime than less stereotypically Black faces.

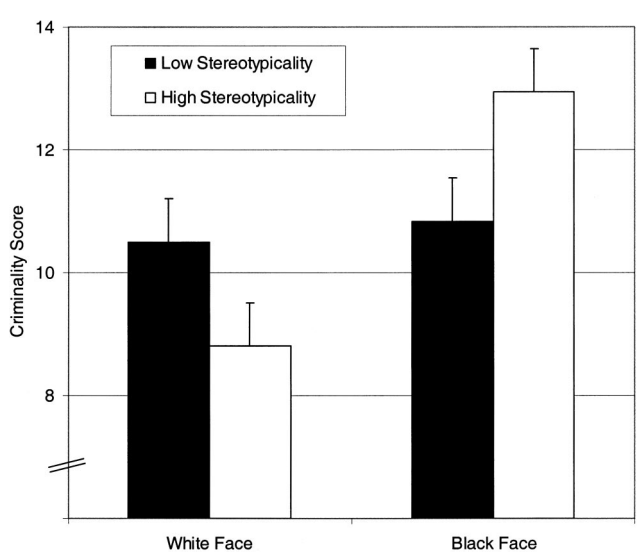

Figure 7. Mean criminality score of faces as a function of race and stereotypicality (Study 5). Error bars represent the average standard error for each condition.

General Discussion

Across five studies, we have shown that bidirectional associations between social groups and concepts can guide how people process stimuli in their visual environment. We found remarkably consistent support for both visual tuning and bidirectionality using three different paradigms that incorporated three different types of participant judgments as well as both image and word stimuli, both student and police officer participant populations, both positive and negative concepts, and both explicit and implicit measures. Specifically, we found that activating stereotypic associations caused participants to detect relevant stimuli at a lower perceptual threshold than irrelevant stimuli (Study 1) and to direct visual attention toward relevant stimuli and away from irrelevant stimuli (Studies 2–4). Furthermore, not only did we demonstrate that social group members bring to mind the concepts with which those social groups are associated (Study 1), we demonstrated that concepts bring to mind the social groups with which those concepts are associated (Studies 2–4). Such effects appear to be related to how strongly a stimulus is thought to represent the social group or concept brought to mind (Study 5).

Our results are consistent with the most recent research findings on stereotypic associations between Black Americans and crime. For instance, Payne and colleagues (Payne, 2001; Payne et al., 2002) found that exposure to Black faces facilitated the categorization of crime-relevant objects. Similarly, Correll and colleagues (2002) found, using a videogame simulation, that participants shot armed Black targets more quickly than armed White targets, irrespective of individual differences in racial attitudes (Correll et al., 2002). Such findings further underscore the strong associational links between Black Americans and crime.

Our research expands previous stereotyping research by more explicitly considering bidirectionality and thus raises new questions about the operation and consequences of stereotypic associations. For example, what determines whether an association will be bidirectional? As discussed earlier, we suspect that concept specificity is one important moderating condition for bidirectionality. Certain concepts may be so tightly coupled with a specific social group that these concepts have become, in a sense, hijacked by that group. Indeed, the social group functions as the prototypical embodiment of these concepts. Concepts for which Black Americans serve as the prototype—such as crime, jazz, basketball, and ghetto—are likely to operate bidirectionally, whereas concepts

[4] A median split was used for ease of presentation. We obtained the same pattern of results when we conducted a regression analysis.

for which there is no specific group prototype—such as aggressive, musical, athletic, and poor—are less likely to operate bidirectionally.

Although not a focus of the current research, situational specificity might also determine the likelihood that a concept will bring to mind a particular social group. For example, when perceivers are required to perform a task that increases the saliency of a particular social group, even concepts that are not attached to any one, prototypical social group may automatically activate the social group that is momentarily salient. For example, aggressive, musical, athletic, and poor are concept primes that may be more likely to activate the Black racial category when the perceivers' subsequent goal is to categorize faces as Black or as White than when the perceivers' goal is race irrelevant (e.g., Kawakami & Dovidio, 2001; Kawakami et al., 2000). Future research studies that systematically manipulate both concept and situational specificity are needed to explore these possibilities more fully.

We also expand previous research on Black Americans and crime by tracing the manner in which such associations can influence critical aspects of visual processing. Specifically, we have shown that activating both positive and negative concepts associated with Blacks (i.e., basketball and crime, respectively) enhances the attentional capture of Black faces (Studies 2–4). We believe that the dot-probe studies introduced here are significant because the results have the potential to increase understanding of attentional selection processes. These results suggest, for example, that stereotypic associations help people to respond to their environment by rendering certain social groups and objects especially relevant and thus worthy of attention. Through these studies we begin to outline the conditions under which members of certain social groups are especially likely to be the objects of gaze, and we examine such visual practices across different participant populations.

In addition, we found that exposure to Black faces allows perceivers to detect crime-relevant objects with less information than exposure to White faces (Study 1). Although detecting objects in an impoverished context is an important skill that numerous vision scientists have sought to understand, researchers have not examined the influence of stereotypic associations on visual practices of this type. We found that Black faces clearly facilitated object detection even in a situation where participants were not required to sort the objects into predetermined categories in a forced-choice format. When participants were given no direction at all on what the object would be, Black faces enhanced their ability to accurately detect degraded crime objects whereas White faces did not. The open-ended format of Study 1 certainly reveals the power of the Black–crime association.

Nevertheless, the design of Study 1 does not permit a precise determination of how such facilitation effects were accomplished. One possibility is that the Black face primes inspired a more detailed, careful visual analysis of crime-relevant objects and that perceivers dedicated more processing resources for this purpose (e.g., see MacLeod et al., 1986). Another possibility is that Black faces inspired less detailed, careful visual inspection of crime-relevant objects rather than more (e.g., see von Hippel, Jonides, Hilton, & Narayan, 1993). Because Black faces are so tightly tied to criminality, exposure to Black faces may have led perceivers to process crime-relevant objects faster but less thoroughly. Participants may have less of a need to inspect crime-relevant objects carefully because they already have some sense of what those objects are. According to this view, exposing perceivers to Black male faces should free up processing resources rather than tax those resources. Distinguishing between these alternatives is an important challenge for future research.

Relatedly, because our results cannot be subjected to signal-detection analysis in any straightforward manner, we are unable to claim with certainty that the Black face primes altered participants' ability to see crime-relevant objects rather than simply rendering participants more confident at identifying indistinct objects as crime relevant. Notably, this latter possibility would require participants to somehow shift their confidence levels beneath awareness (given that the face primes were subliminal) and in opposite directions for Black and White faces. In addition, if the results were simply due to shifts in confidence levels, one might expect the error rates across conditions to fluctuate—which they did not.

A central theme in perception research concerns the mechanisms that give rise to everyday visual experience in a world that exposes perceivers to more sensory information than they have the capacity to process thoroughly. Thus far, perception researchers have approached this issue by examining both the properties of stimulus inputs and the computational properties of the brain. In many models of perception, conceptual knowledge is thought to modulate visual processes in important ways. However, despite the importance of social knowledge in everyday interactions, the role of social knowledge in visual processing is rarely discussed. We have argued that visual analysis is, in part, socially driven. Stereotypic associations, in particular, have the capacity to critically alter visual experience.

Throughout this article, we have suggested that the effects of stereotypic associations on visual perception and attention could be of great practical significance. Indeed, given the perceptual threshold effects reported here, police officers may face elevated levels of danger in the presence of White armed suspects in comparison with Black armed suspects. For example, if police officers have a delayed response to White suspects with guns or knives, these officers may be more likely to get hurt, shot, or killed when confronting White armed suspects in comparison with Black armed suspects. In contrast, unarmed, innocent Blacks may easily become the targets of intense visual surveillance by both police officers and the lay public. With their eyes, perceivers may tie individual Black targets to a group-based suspicion—and sadly, Black people who appear highly stereotypically Black may be the most likely of all to feel the tug. Such processes could indicate that racial profiling may be rooted in more fundamental perceptual processes than previously recognized.

The studies presented here might have implications for the durability of stereotypic associations. Numerous factors may conspire to maintain the strength of such associations. For one thing, to the extent that these associations are bidirectional, they can be activated even in the absence of a social category member. Concepts that are represented by a prototypical racial category have the power to conjure their own racialized subjects. Indeed, we believe that notions of race are so powerful because they can operate through ostensibly race-neutral concepts (such as crime).

Additionally, although not addressed in the work presented here, the motivation of perceivers to actively resist stereotypic associations may differ on the basis of whether these associations are triggered by social group members or concepts. For instance, many

people may be motivated to actively resist thoughts of criminality in the presence of a Black American trigger (e.g., see Dunton & Fazio, 1997; Plant & Devine, 1998). Activations of this type may even be considered a personal failing (to the extent that perceivers are aware of them and wish to be egalitarian). In contrast, people may be less motivated to resist thoughts of Black Americans in the presence of a crime trigger. Far from a personal failing, the activation of such thoughts may be experienced as a natural response, given the high proportion of Blacks convicted of violent crimes in the United States (Banks, 2001; Blumstein, 1993; Cole, 1999; Kennedy, 1997). In fact, we have preliminary evidence with police officer participants suggesting that the motivation to resist stereotypic associations may depend on the triggering stimulus (Eberhardt & Goff, 2004). We found that police officers are less troubled by the possibility of crime triggering thoughts of Black Americans than by the possibility of Black Americans triggering thoughts of crime. We suspect that this asymmetry is present in American society more generally. Egalitarian opposition to racial stereotyping strongly condemns linking Black people to crime, but not linking crime to Black people. Thus, opposition to stereotyping tends to condemn one aspect of the association, even as it exempts the other.

Finally, visual practices may not simply reflect race-based associations; visual practices may work to sustain these associations as well. Visual processing patterns may provide ample opportunities for perceivers to access race–crime associations, as well as to rehearse, strengthen, and supplement those associations. In this way, seeing could be understood as an action or a practice that reinscribes racial meaning onto visual stimuli. The face-recognition memory results of Study 4 are consistent with such an interpretation. Activation of the crime concept not only led police officers to attend to a Black face but also led them to misremember the Black face as more stereotypical (i.e., representative) of the Black racial category than it actually was. Thus, the association between blackness and criminality was not only triggered, it was magnified.

It is important to note that although visual processes may reinforce stereotypic associations, the associations themselves are the consequences of widely shared cultural understandings and social patterns. As William James stated, attention "creates no idea." Because visual processes are grounded in cultural understandings, as these understandings change, the consequences of visual processes will as well. New associations may render different aspects of the visual environment relevant and expose perceivers to a different world from the one they currently have the capacity to see.

References

Allport, G. W., & Postman, L. J. (1947). *The psychology of rumor*. New York: Russell & Russell.

Anderson, S., & Klatzky, R. (1987). Traits and social stereotypes: Levels of categorization in person perception. *Journal of Personality and Social Psychology, 53,* 235–246.

Asch, S. E., & Ebenholtz, S. M. (1962). The principle of associative memory. *Proceedings of the American Philosophical Society, 106,* 135–163.

Banks, R. R. (2001). Race-based suspect selection and colorblind equal protection doctrine and discourse. *UCLA Law Review, 48,* 1075–1124.

Bargh, J. A., & Chartrand, T. L. (2000). The mind in the middle: A practical guide to priming and automaticity research. In H. T. Reis & C. M. Judd (Eds.), *Handbook of research methods in social and personality psychology* (pp. 253–285). New York: Cambridge University Press.

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

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

Blair, I. V., Judd, C. M., Sadler, M. S., & Jenkins, C. (2002). The role of Afrocentric features in person perception: Judging by features and categories. *Journal of Personality and Social Psychology, 83,* 5–25.

Blumstein, A. (1993). Racial disproportionality of U.S. prison populations revisited. *University of Colorado Law Review, 64,* 743–760.

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 and Emotion, 12,* 737–753.

Bradley, B. P., Mogg, K., Millar, N., Bonham-Carter, C., Fergusson, E., Jenkins, J., & Parr, M. (1997). Attentional biases for emotional faces. *Cognition and Emotion, 11,* 25–42.

Brewer, M. B., Dull, V., & Lui, L. (1981). Perceptions of the elderly: Stereotypes as prototypes. *Journal of Personality and Social Psychology, 41,* 656–670.

Chen, M., & Bargh, J. A. (1997). Nonconscious behavioral confirmation processes: The self-fulfilling consequences of automatic stereotype activation. *Journal of Experimental Social Psychology, 33,* 541–560.

Cole, D. (1999). *No equal justice: Race and class in the American criminal justice system.* New York: New Press.

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

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

Dolan, R. J., Fink, G. R., Rolls, E., Booth, M., Holmes, A., Frackowiak, R. S. J., & Friston, K. J. (1997). How the brain learns to see objects and faces in an impoverished context. *Nature, 389,* 596–599.

Dovidio, J., Evans, N., & Tyler, R. (1986). Racial stereotypes: The contents of their cognitive representation. *Journal of Experimental Social Psychology, 22,* 22–37.

Dovidio, J. F., Kawakami, K., Johnson, K., Johnson, C., & Howard, A. (1997). On the nature of prejudice: Automatic and controlled processes. *Journal of Experimental Social Psychology, 33,* 510–540.

Duncan, B. L. (1976). Differential social perception and attribution of intergroup violence: Testing the lower limits of stereotyping blacks. *Journal of Personality and Social Psychology, 34,* 590–598.

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

Eberhardt, J. L., Dasgupta, N., & Banaszynski, T. L. (2003). Believing is seeing: The effects of racial labels and implicit beliefs on face perception. *Personality and Social Psychology Bulletin, 29,* 360–370.

Eberhardt, J. L., & Goff, P. A. (2004). [The morality of associating race and crime]. Unpublished raw data.

Ellison, R. (1952). *Invisible man*. New York: Random House.

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.

Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology, 50,* 229–238.

Fazio, R. H., Williams, C. J., & Powell, M. C. (2000). Measuring associative strength: Category-item associations and their activation from memory. *Political Psychology, 21,* 7–25.

Gaertner, S. L., & McLaughlin, J. P. (1983). Racial stereotypes: Associations and ascriptions of positive and negative characteristics. *Social Psychology Quarterly, 46,* 23–30.

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.

Gilboa-Schechtman, E., Foa, E. B., & Amir, N. (1999). Attentional biases for facial expressions in social phobia: The face-in-the-crowd paradigm. *Cognition and Emotion, 13,* 305–318.

Goldstein, E. B. (1999). *Sensation and perception* (5th ed.). Pacific Grove, CA: Brooks/Cole.

Gotlib, I. H., McLachlan, A. L., & Katz, A. N. (1988). Biases in visual attention in depressed and nondepressed individuals. *Cognition and Emotion, 2,* 185–200.

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

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

Grill-Spector, K., Kushnir, T., Hendler, T., & Malach, R. (2000). The dynamics of object-selective activation correlate with recognition performance in humans. *Nature Neuroscience, 3,* 837–843.

Horowitz, L. M., Norman, S. A., & Day, R. S. (1966). Availability and associative memory. *Psychological Review, 73,* 1–15.

Horowitz, L. M., & Prytulak, L. S. (1969). Redintegrative memory. *Psychological Review, 76,* 519–531.

Hunt, R. G. (1959). Meaningfulness and articulation of stimulus and response in paired-associate learning and stimulus recall. *Journal of Experimental Psychology, 57,* 262–267.

James, W. (1950). *The principles of psychology.* New York: Dover. (Original work published 1890)

Jantz, E. M., & Underwood, B. J. (1958). R-S learning as a function of meaningfulness and degree of S-R learning. *Journal of Experimental Psychology, 56,* 174–179.

Johnson, W. A., & Dark, V. J. (1986). Selective attention. In M. R. Rosenzweig & L. W. Porter (Eds.), *Annual review of psychology* (pp. 43–75). Palo Alto, CA: Annual Reviews.

Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1999). Increased activity in the human visual cortex during directed attention in the absence of visual stimulation. *Neuron, 22,* 751–761.

Kawakami, K., Dion, K. L., & Dovidio, J. (1998). Racial prejudice and stereotype activation. *Personality and Social Psychology Bulletin, 24,* 407–416.

Kawakami, K., & Dovidio, J. F. (2001). The reliability of implicit stereotyping. *Personality and Social Psychology Bulletin, 27,* 212–225.

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

Kennedy, R. (1997). *Race, crime, and the law.* New York: Pantheon Books.

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

Livingston, R. W. (2001). What you see is what you get: Systematic variability in perceptual-based social judgment. *Personality and Social Psychology Bulletin, 27,* 1086–1096.

MacLeod, C., Mathews, A., & Tata, P. (1986). Attentional bias in emotional disorders. *Journal of Abnormal Psychology, 95,* 15–20.

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

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

Maddox, K. B., & Gray, S. A. (2001). Cognitive representations of Black Americans: Reexploring the role of skin tone. *Personality and Social Psychology Bulletin, 28,* 250–259.

Maidenberg, E., Chen, E., Craske, M., Bohn, P., & Bystritsky, A. (1996). Specificity of attentional bias in panic disorder and social phobia. *Journal of Anxiety Disorders, 10,* 529–541.

Mansell, W., Clark, D. M., Ehlers, A., & Chen, Y.-P. (1999). Social anxiety and attention away from emotional faces. *Cognition and Emotion, 13,* 673–690.

Martin, M., Williams, R. M., & Clark, D. M. (1991). Does anxiety lead to selective processing of threat-related information? *Behaviour Research and Therapy, 29,* 147–160.

Mathews, A., & MacLeod, C. (1985). Selective processing of threat cues in anxiety states. *Behaviour Research and Therapy, 23,* 563–569.

McConahay, J. B. (1986). Modern racism, ambivalence, and the Modern Racism Scale. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 91–125). Orlando, FL: Academic Press.

Mogg, K., & Bradley, B. P. (1999). Some methodological issues in assessing attentional biases for threatening faces in anxiety: A replication study using a modified version of the probe detection task. *Behaviour Research and Therapy, 37,* 595–604.

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

Neely, J. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General, 106,* 226–254.

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,* 1–12.

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. (2003, February). *Cognitive control and the experience of misperceiving a weapon.* Paper presented at the Fourth Annual Meeting of the Society for Personality and Social Psychology, Los Angeles, CA.

Perdue, C., & Gurtman, M. (1990). Evidence for the automaticity of ageism. *Journal of Experimental Social Psychology, 28,* 199–216.

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

Ress, D., Backus, B. T., & Heeger, D. J. (2000). Activity in primary visual cortex predicts performance in a visual detection task. *Nature Neuroscience, 3,* 940–945.

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.

von Hippel, W., Jonides, J., Hilton, J. L., & Narayan, S. (1993). Inhibitory effect of schematic processing on perceptual encoding. *Journal of Personality and Social Psychology, 64,* 921–935.

von Hippel, W., Sekaquaptewa, D., & Vargas, P. (1995). On the role of

encoding processes in stereotype maintenance. *Advances in Experimental Social Psychology, 27,* 177–254.

Williams, M. J., & Eberhardt, J. L. (2004). *Racial conceptions: Measuring individual differences in lay theories of race and motivation to cross racial boundaries*. Manuscript in preparation.

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

Yantis, S., & Egeth, H. E. (1999). On the distinction between visual salience and stimulus-driven attentional capture. *Journal of Experimental Social Psychology, 25,* 661–676.

Received July 10, 2003
Revision received August 14, 2004
Accepted August 14, 2004 ■

AMERICAN PSYCHOLOGICAL ASSOCIATION
SUBSCRIPTION CLAIMS INFORMATION

Today's Date:_____

We provide this form to assist members, institutions, and nonmember individuals with any subscription problems. With the appropriate information we can begin a resolution. If you use the services of an agent, please do **NOT** duplicate claims through them and directly to us. **PLEASE PRINT CLEARLY AND IN INK IF POSSIBLE.**

PRINT FULL NAME OR KEY NAME OF INSTITUTION

ADDRESS

CITY STATE/COUNTRY ZIP

YOUR NAME AND PHONE NUMBER

MEMBER OR CUSTOMER NUMBER (MAY BE FOUND ON ANY PAST ISSUE LABEL)

DATE YOUR ORDER WAS MAILED (OR PHONED)

____PREPAID ____CHECK ____CHARGE
CHECK/CARD CLEARED DATE:_____

(If possible, send a copy, front and back, of your cancelled check to help us in our research of your claim.)

ISSUES: ____ MISSING ____ DAMAGED

TITLE	VOLUME OR YEAR	NUMBER OR MONTH

Thank you. Once a claim is received and resolved, delivery of replacement issues routinely takes 4–6 weeks.

(TO BE FILLED OUT BY APA STAFF)

DATE RECEIVED: _____ **DATE OF ACTION:** _____
ACTION TAKEN: _____ **INV. NO. & DATE:** _____
STAFF NAME: _____ **LABEL NO. & DATE:** _____

Send this form to APA Subscription Claims, 750 First Street, NE, Washington, DC 20002-4242

PLEASE DO NOT REMOVE. A PHOTOCOPY MAY BE USED.