UC Davis

UC Davis Previously Published Works

Title

No face is an island: How implicit bias operates in social scenes

Permalink

<https://escholarship.org/uc/item/1bf0j1sv>

Journal

Journal of Experimental Social Psychology, 49(2)

ISSN

0022-1031

Authors

Soderberg, CK

Sherman, JW

Publication Date

2013-03-01

DOI

10.1016/j.jesp.2012.11.001

Peer reviewed

Reports

No face is an island: How implicit bias operates in social scenes

Courtney K. Soderberg*, Jeffrey W. Sherman*

Department of Psychology, University of California, One Shields Avenue, Davis, CA 95616, USA

HIGHLIGHTS

- We examined how viewing diverse and homogenous social contexts affects implicit bias.
- Racially diverse contexts decrease bias and homogenous contexts increase bias.
- Decreases in bias generalized to situations in which targets were seen in isolation.
- Diverse contexts did not decrease bias toward targets never seen in diverse contexts.
- Quad modeling revealed effects were related to changes in automatic associations.

ARTICLE INFO

Article history:

Received 20 February 2012

Available online 19 November 2012

Keywords:

Implicit bias
Social context
Process model

ABSTRACT

Social psychologists have mainly studied implicit attitudes toward faces presented one at a time, whereas, in real life, we often encounter people in the presence of others. These surrounding individuals may alter attitudes toward the focal target of attention. We employed a flanker-IAT task and found that, when black and white targets were presented in racially diverse contexts, bias was decreased. This decrease in bias occurred even when targets previously seen in diverse contexts were presented on their own, suggesting context-free evaluations of the targets had been formed. Experiment 2 showed that the effect of diverse contexts does not affect bias toward a racial category as a whole, but only the specific targets previously seen in the diverse contexts. Quad model analysis (Sherman et al., 2008) revealed that these effects were related to changes in automatic evaluations, and not to changes in inhibition. Implications for implicit bias change and prejudice reduction strategies are discussed.

© 2012 Elsevier Inc. All rights reserved.

Introduction

Over the past 15 years, social psychologists have documented a multitude of factors that influence implicit racial bias. Many aspects of stimuli, such as the specific faces presented (e.g., Dasgupta & Greenwald, 2001) or the physical environment in which the faces are placed (e.g., Wittenbrink, Judd, & Park, 2001) have been shown to alter attitudes toward racial groups. Although this research has advanced the understanding of implicit attitude malleability, one central aspect of the social world has yet to be examined: namely, the presence of multiple social targets. At present, social psychologists have tested the operation of implicit bias only in situations in which perceivers observe a single target. However, whether milling around in a crowd, talking to friends, or sitting next to a stranger on a subway, perceivers often encounter people in the presence of others. Although attention may be focused on one individual, the other people present in the visual scene may influence impressions of the focal target. The ubiquity of this type of setting suggests that, in order to better understand how implicit

bias operates in complex social environments, researchers should investigate whether and how the presence of others in the visual scene affects bias toward a focal individual. In the present research, we were particularly interested in scenes involving multiple targets that are either the same race as the focal target or a different race. Specifically, we investigated the influence of racially homogenous versus diverse contexts on implicit racial bias.

Although there is no research that has directly tested how the presence of others in the visual scene might affect bias toward a target individual, it is possible to derive two competing hypotheses based on past theory and research. On the one hand, prejudice toward a target could decrease when the stimulus is placed in a racially diverse as opposed to a racially homogenous context. Research on the effects of context on implicit attitudes has overwhelmingly found assimilation effects—targets being evaluated in line with the valence of the context (Allen, Sherman, & Klauer, 2010; Barden, Maddux, Petty, & Brewer, 2004; Wittenbrink et al., 2001). In the present study, such an assimilation effect would lead black targets to be evaluated more favorably when surrounded by white people, for example. Such effects may occur for various reasons, including context-based evaluative conditioning, changes in category activation or inhibition processes, or subtyping. These possibilities will be examined in detail in the General discussion.

* Corresponding authors. Fax: +1 530 752 2087.

E-mail addresses: cksoderberg@ucdavis.edu (C.K. Soderberg),
jsherman@ucdavis.edu (J.W. Sherman).

On the other hand, prejudice toward a target might increase when the stimulus is placed in a racially diverse as opposed to a racially homogenous context. Numerical distinctiveness theory (McGuire, McGuire, Child, & Fujioka, 1978) proposes that categories that are numerically distinct in interpersonal settings draw attention and become more salient because attention is drawn to novel stimuli, a process that leads targets to be contrasted away from the valence of the social context. In fact, research on category activation and explicit stereotyping has shown primarily contrast effects in interpersonal contexts (e.g., McGuire et al., 1978; Oakes & Turner, 1986; Stroessner, 1996; Zarate & Smith, 1990). It is worth noting that the research described above demonstrating assimilation of implicit prejudice to context relied solely on presentations of person targets in front of environmental background contexts. However, unlike environmental background contexts, other humans provide a context that is perceptually and conceptually similar to target persons. This overlap between target and context may enhance comparison processes that produce contrast rather than assimilation effects in evaluation and judgment (e.g., Herr, 1986; Herr, Sherman, & Fazio, 1983). In the present study, such a contrast effect would lead black targets to be evaluated less favorably when surrounded by white people, for example.

Experiment 1

In Experiment 1, we created a version of the IAT (Greenwald, McGhee, & Schwartz, 1998) in which three stimuli appeared on the screen at once. We used this task to test how the level of implicit bias elicited by racially diverse versus homogenous contexts compared to the level elicited by individuals viewed in isolation.

Method

Subjects

Fifty-nine non-black University of California, Davis undergraduates participated in the study for partial course credit.

Stimuli

We used a standard evaluative IAT to measure implicit prejudice by assessing the ease with which participants were able to associate white and black faces with pleasant and unpleasant words. In addition, we incorporated aspects of an Eriksen-Flanker task (Eriksen & Eriksen, 1974)—in which a central target is surrounded by either congruent or incongruent flankers—into the IAT in order to determine how surrounding individuals affect prejudice toward a focal target. This created three different types of IAT trials: unflanked trials, in which one stimulus was presented on any given trial; racially homogenous flanked trials, in which three stimuli appeared on the screen simultaneously to present homogenous contexts (e.g., a white target surrounded by white flankers); and racially diverse flanked trials, in which three stimuli appeared on the screen simultaneously to present racially diverse contexts (e.g., a black target surrounded by white flankers) (Fig. 1). The same set of stimuli was used to create the unflanked and flanked trials.

Procedure

Participants were randomly assigned to one of three conditions: a control condition in which they completed an evaluative IAT with only unflanked trials; a racially homogenous condition that included both unflanked and homogenous-flanked IAT trials; and a racially diverse condition that included both unflanked and diverse-flanked IAT trials. For flanked trials, participants were instructed to ignore the surrounding stimuli and respond to the central target stimulus. Unflanked trials were included in the homogenous and diverse conditions to determine whether the effects of the flankers extended to trials in which targets were not flanked. For the diverse and homogenous conditions, flanked and unflanked trials were randomly ordered, so that participants could not anticipate trial type.

Fig. 1. A) Example of a racially diverse black flanked stimulus B) Example of a racially homogenous black flanked stimulus.

Results

IAT scores were calculated using the improved scoring algorithm (Greenwald, Nosek, & Banaji, 2003). For subjects in the homogenous and diverse conditions the two trial types were analyzed separately to produce one IAT score for flanked trials and another for unflanked trials. This allowed us to test for the effect of social context on implicit bias by comparing IAT scores from flanked trials to the control condition. It also allowed us to test whether the effects of viewing social contexts extended to situations in which individuals were later seen in isolation by comparing IAT scores from the unflanked trials of the homogenous and diverse conditions to the control condition.

To test whether social context affects bias toward target individuals, we conducted a univariate ANOVA on the unflanked control condition and the flanked trials from both the homogenous and diverse conditions, which demonstrated a significant effect, $F(2, 56) = 9.99, p = .00, \eta^2 = .26$. Follow-up pairwise comparisons using Fisher's LSD (Cohen, 2001) showed that, as compared to the control condition ($M = .54, SD = .36$), homogenous contexts raised anti-black/pro-white bias, ($M = .82, SD = .23$), $t(56) = -2.43, p = .02, d = .94$, whereas diverse contexts lowered anti-black/pro-white bias, ($M = .31, SD = .44$), $t(56) = 2.07, p = .04, d = .59$ (Fig. 2).¹

To test whether the effect of social context on bias extended to trials in which individuals previously seen in those contexts are now viewed in isolation, we conducted a univariate ANOVA on the unflanked control condition and the unflanked trials from both the homogenous and diverse conditions, which demonstrated a significant effect, $F(2, 56) = 4.57, p = .02, \eta^2 = .14$. Follow-up pairwise comparisons using Fisher's LSD to compare the unflanked trials from the homogenous and diverse conditions to the control condition revealed that, although the control condition ($M = .54, SD = .36$) did not differ from the unflanked trials in the homogenous condition ($M = .60, SD = .29$), $p = .64, d = .16$, it did significantly differ from the unflanked trials in the diverse condition ($M = .29, SD = .37$), $p = .02, d = .72$. These results indicate that seeing diverse social contexts not only

¹ For flanked trials, both words and faces were flanked in order to allow for Quad model analysis. Because the type of focal stimulus in the trial (words versus faces) made no difference in the analyses, we collapsed across type of stimulus in the analyses reported here, as is standard in IAT analysis.

Fig. 2. IAT scores, Experiment 1.

lowered bias on flanked trials, but also lowered bias on subsequently encountered unflanked trials, showing that the effects of the diverse contexts extended to situations in which only a single target was observed.²

Discussion

Experiment 1 showed that diverse flanked trials reduced pro-white bias and that homogenous flankers increased pro-white bias on the IAT. Additionally, Experiment 1 showed that the decrease in bias produced by diverse flanked trials extended to unflanked trials whose stimuli had previously been seen in flanked arrays. This suggests that the decrease in bias may be maintained, for a particular target, beyond the particular time and context in which that target is initially encountered. However, due to the fact that the same stimuli were used for both the flanked and unflanked trials, it is unclear from these results whether the flanked trials were altering perceptions of only the individual stimuli presented or perceptions of the categories as a whole. We designed Experiment 2 to test this question.

Experiment 2

In Experiment 2, we examined whether racially diverse contexts change only the perceptions of the particular targets encountered in those contexts or whether the effects generalize to perceptions of the target's category as a whole. In other words, does seeing a black face surrounded by two white faces change perceptions of the specific black face or does it change perceptions of the category of black people as a whole? To investigate this question, diverse flanked trials were interspersed with unflanked trials containing stimuli that had either previously been encountered in a flanked trial or had never been seen during flanked trials. If the diverse contexts only alter perceptions of specific targets encountered in diverse contexts, then only

stimuli that have previously been seen in those arrays should elicit lower bias when they are seen on their own. In other words, the carry-over effect should only hold for those specific targets. By contrast, if the diverse contexts alter the perceptions of the category as a whole, then new stimuli that have never been seen in flanked arrays should also elicit lower bias when presented alone. We also expect that all diverse flanked trials should show less bias than a control condition, replicating the results of Experiment 1.

Method

Subjects

Seventy-two non-black University of California, Davis undergraduates participated for partial course credit.

Procedure

Participants were randomly assigned to one of three conditions: a *control* condition in which they completed an evaluative IAT with only unflanked trials; a *familiar* condition in which they completed both unflanked and diverse flanked trials, with both types of trials created using the same stimuli (as in Experiment 1); or a *novel* condition in which they completed both unflanked and diverse flanked trials, with unflanked and flanked trials created using two different sets of stimuli.

Results

IAT scores were calculated using the improved scoring algorithm (Greenwald et al., 2003). For subjects in the familiar and novel conditions who completed both flanked and unflanked IAT trials, the two trial types were analyzed separately, as in Experiment 1.

In order to test whether the effect found in Experiment 1 was replicated in Experiment 2, a planned contrast was done to examine whether the control condition ($M = .69$, $SD = .37$) showed more bias than the flanked trials from the familiar ($M = .36$, $SD = .40$) and novel conditions ($M = .48$, $SD = .41$). The comparison was significant, $t(70) = 2.78$, $p = .01$, replicating the findings of Experiment 1.

Our central question in Experiment 2 was whether the decrease in prejudice brought about by viewing diverse contexts only extended to familiar targets viewed in isolation or whether it also extended to novel targets viewed in isolation. In order to test this, we conducted a univariate ANOVA on the control and unflanked trials from both the familiar and novel conditions, which was marginally significant, $F(2,70) = 2.48$, $p = .09$, $\eta^2 = .07$ (Fig. 3). Planned comparisons showed that the unflanked trials in the familiar condition ($M = .45$, $SD = .41$) showed significantly less bias than the control condition ($M = .69$, $SD = .37$), $t(70) = 2.22$, $p = .03$, $d = .63$, replicating the results from Experiment 1. In contrast, the unflanked trials from the novel condition ($M = .58$, $SD = .34$) did not show a significantly different level of bias than the control condition, $t(70) = 1.01$, $p = .32$, $d = .31$, indicating that the decrease in bias caused by diverse contexts does not generalize to the categories as a whole, but only to faces that have previously been observed in flanked arrays (see Fig. 3).³

Identifying the mechanism with modeling

Consistent with previous research, these two experiments showed that implicit evaluations of target persons shifted toward the evaluative implications of the context in which they were placed. There are at least two mechanisms that may have produced these assimilation effects.

² To show that the decrease in bias caused by viewing targets in diverse contexts carried over to situations in which the same stimuli were later seen in isolation, unflanked trials whose stimuli had not yet been seen in a flanked array were excluded from this analysis. This resulted in the exclusion of 7% of the unflanked trials in the *homogenous* condition and 8% in the *diverse* condition. Inclusion of these trials did not alter the results of any analyses.

³ As in Experiment 1, unflanked trials in the *familiar* conditions whose stimuli had not yet been seen in a flanked array were excluded from analysis. This resulted in the exclusion of 10% of the unflanked trials from the *familiar* condition. Inclusion of these trials did not change the results of any analyses.

Fig. 3. IAT scores, Experiment 2.

First, diverse flankers may alter perceivers' evaluations of the target faces. This may occur if diverse contexts decrease the extent to which a target face is racially categorized or activate counter-stereotypical group subtypes. Alternatively, target evaluations may change if the evaluative implications of the diverse contexts may bleed over to influence the evaluations of the target faces, as in evaluative conditioning (e.g., Jones, Fazio, & Olson, 2009; Walther, 2002).

A second explanation for the decrease in bias in diverse contexts is that the diverse flankers cue perceivers to control their responding. In this case, perception of the target is unaffected, but the context prompts people to inhibit biased responses to the target. When people detect a discrepancy between how they are responding and how they believe they should respond, they attempt to inhibit their behavior and develop environmental cues for control that signal that preemptive behavioral

inhibition is necessary (Monteith, Ashburn-Nardo, Voils, & Czopp, 2002). Indeed, when black targets are presented in positive contexts (e.g., church), they evoke less implicit bias along with more effective inhibition of biased associations than when they are presented in negative contexts (e.g., jail; Allen et al., 2010). Similarly, it may be that the surrounding individuals in diverse contexts serve as cues to participants to control their responding, leading to a decrease in bias. Thus, the decrease in bias observed in our studies could reflect either a change in the evaluations of the central targets or an increase in participants' inhibition of biased responding.

In order to shed light on which of these alternative mechanisms can account for our results, we applied the Quadruple Process model (Quad model; Conrey, Sherman, Gawronski, Hugenberg, & Groom, 2005; Sherman et al., 2008) to the IAT data from Studies 1 and 2. The Quad model is a multinomial model designed to estimate the independent contributions of multiple processes to responses on implicit measures of bias. According to the model, performance on implicit measures reflects four distinct processes: Activation of Associations (AC)—the degree to which associations are activated when encountering a stimulus; Detection (D)—the ability to detect the correct response; Overcoming Bias (OB)—a self-regulatory process that overcomes the influence of associations when they conflict with correct responses; and Guessing (G)—general response tendencies that occur when associations are not activated and correct responses cannot be determined.

The structure of the model is depicted as a processing tree in Fig. 4. Each path in the tree represents a compound probability (e.g., $AC \times D \times OB$) and predicts a specific response (i.e., correct or incorrect). The sum of all probabilities associated with a response is the total probability of that response. For example, when presented with a black face in the incompatible block of a black–white IAT, the probability of a correct response is $[AC \times D \times OB] + [(1 - AC) \times D] + [(1 - AC) \times (1 - D) \times (1 - G)]$. This equation sums the three possible paths by which a correct answer will occur. The first part of the equation, $AC \times D \times OB$, is the likelihood that the association is activated, that the correct response can be detected, and that the association is overcome. The second part of the equation, $(1 - AC) \times D$, is the likelihood that the association is not activated and that the correct response can be determined. Finally, $(1 - AC) \times (1 - D) \times (1 - G)$ is the likelihood that the association is not activated, that the correct response cannot be determined, and that

Fig. 4. The Quadruple Process model (Quad model). Each path represents a likelihood. Parameters with lines leading to them are conditional upon all preceding parameters. The table on the right side of the figure depicts correct (✓) and incorrect (✗) responses as a function of process pattern and trial type. In this particular figure, the guessing bias refers to guessing in line with activated associations.

Table 1
Parameter estimates for black–white IAT from Experiment 1.

Parameter	Estimate		
	Unflanked control condition trials	Homogenous flanked trials	Diverse flanked trials
AC			
Black-bad	0.08 (0.02)	0.29 (0.04)	0.01 (0.03)
White-good	0.12 (0.02)	0.18 (0.03)	0.04 (0.03)
OB	0.80 (0.14)	0.78 (0.10)	1.00 (0.99)
D	0.90 (0.01)	0.91 (0.02)	0.83 (0.02)
G	0.57 (0.05)	0.46 (0.10)	0.53 (0.05)

Notes. AC = Activation of Associations, D = Detection, OB = Overcoming Bias, G = Guessing.

the participant guesses correctly. The sum of these probabilities is the total probability of a correct response for the item.

The four parameter values are changed through maximum likelihood estimation to maximize fit between the model and observed responses. The parameter values resulting from this procedure represent the levels of the corresponding processes during task performance. The Quad model and the construct validity of its parameters have been extensively validated in previous research (Conrey et al., 2005; Sherman et al., 2008).

For each condition, we estimated one AC parameter that measured the association between black faces and unpleasant words, one AC parameter that measured the association between white faces and pleasant words, one OB, one D, and one G parameter. If the decrease in bias seen in diverse contexts is due to a change in the evaluation of the stimuli, then we would expect lower levels of AC in diverse flanked trials than in unflanked trials. In contrast, if the effects are due to self-regulation, then we would expect higher levels of OB in the diverse flanked trials than in the unflanked trials.

All trials from the control condition and the diverse flanked trials from the diverse condition of Experiment 1 were subjected to Quad model analysis. The Quad model fit the data, $G^2(2) = 3.67$, $p = .16$. Results showed that the AC black-bad parameter was significantly lower for the diverse flanked trials than for the control condition unflanked trials, $G^2(1) = 3.94$, $p = .05$. The AC white-good parameter also was significantly lower on the diverse flanked trials than on the control condition unflanked trials, $G^2(1) = 4.96$, $p = .03$. The OB parameter did not differ between the control and diverse conditions, $G^2(1) = 0.38$, $p = .54$ (Table 1). These results show that target faces activated less biased evaluative associations in the diverse condition, indicating that the decrease in prejudice seen in this condition was due to changes in the manner in which targets were automatically evaluated.⁴

To test whether a similar process characterized our results in Experiment 2, we conducted a Quad model analysis of the control condition and the flanked trials from the familiar and novel conditions in this experiment. The flanked trials from the novel and familiar conditions were both diverse flanked trials and showed lower IAT bias than the control condition. As such, for the Quad model analysis of Experiment 2, the flanked trials from the novel and familiar conditions were combined so that the model tested for parameter differences between unflanked control trials and diverse flanked trials, as in the analysis from Experiment 1. The model fit the data, $G^2(2) = 4.66$, $p = .10$. As in Experiment 1, the AC black-bad parameter was significantly lower for the diverse flanked trials than for the unflanked control condition trials, $G^2(1) = 6.31$, $p = .01$. The AC white-good parameter also was significantly lower for the diverse flanked trials than

Table 2
Parameter estimates for black–white IAT from Experiment 2.

Parameter	Estimate	
	Unflanked control condition trials	Diverse flanked trials
AC		
Black-bad	0.08 (0.02)	0.03 (0.01)
White-good	0.10 (0.02)	0.05 (0.01)
OB	0.26 (0.18)	0.00 (0.41)
D	0.91 (0.01)	0.85 (0.01)
G	0.50 (0.05)	0.52 (0.03)

Notes. AC = Activation of Associations, D = Detection, OB = Overcoming Bias, G = Guessing.

for the unflanked control condition trials, $G^2(1) = 7.07$, $p < .01$. The OB parameter did not differ in the two conditions, $G^2(1) = 1.23$, $p = .27$, and was, in fact, numerically lower in the diverse condition (Table 2). These findings replicate the results of Experiment 1, and provide further support for the conclusion that changes in target evaluation rather than changes in self-regulation is the mechanism underlying the decrease in implicit bias observed in diverse settings.⁵

We also tested whether the increase in bias seen in homogenous contexts was related to changes in target evaluations. To this end, we subjected the control condition and the homogenous flanked trials from the homogenous condition from Experiment 1 to Quad model analysis. The Quad model fit the data, $G^2(2) = 3.06$, $p = .22$. Results showed that the AC black-bad parameter was significantly higher for the flanker trials than the control trials, $G^2(1) = 26.55$, $p < .001$. The AC white-good parameter did not differ between conditions, $G^2(1) = 2.33$, $p = .13$. The OB parameter also did not differ between conditions, $G^2(1) = 0.02$, $p = .90$ (Table 1). These results show that the increase in bias in homogenous contexts is related to changes in automatic black-negative associations, again suggesting that the contexts directly affected the evaluations of target faces.

In sum, the results of the Quad model analysis indicate that changes in implicit bias caused by interpersonal contexts are due to changes in the way that targets are evaluated, not to changes in the regulation of evaluative bias. Thus, although both interpersonal and background environmental contexts lead to assimilation effects—the target being evaluated in line with the contexts—the mechanism underlying these effects appears to be different for the two types of context (e.g., Allen et al., 2010).

General discussion

Two studies showed that the presence of other people in a visual scene has a significant impact on implicit bias. In particular, when black and white targets were encountered in racially homogenous contexts, implicit bias increased compared to when they were encountered in isolation. In contrast, when targets were encountered in racially diverse contexts, implicit bias decreased. Additionally, this decrease in bias generalized to situations in which targets that had previously been seen in diverse contexts were subsequently encountered in isolation, suggesting that context-free evaluations of the central targets were formed. However, Experiment 2 revealed that this decrease in bias did not generalize to novel stimuli, indicating that the presence of diverse social contexts changes attitudes toward particular group members but not toward the groups as a whole. Application of the Quad model to the data showed that these effects were related to contextual changes in the evaluations of targets and not to changes in the extent of inhibition triggered by the contexts. These findings are, to our knowledge, the first to show how the presence of other people affects implicit prejudice toward target individuals in social contexts. They demonstrate that social scenes can have a powerful and lasting impact on

⁴ Modeling also showed that the D parameter was lower for diverse flanked trials than for the unflanked control condition trials, which likely reflects the greater difficulty of identifying stimuli on flanked trials. Note that, according to the Quad model's equations, higher D should reduce bias. As such, the observed differences in D cannot account for reduced bias in the diverse condition.

⁵ As in Experiment 1, the D parameter was smaller for diverse flanked trials than for the unflanked control condition trials.

how we perceive individuals across situations. Given that, in everyday life, people frequently encounter more than one individual at a time, studying these more complex social situations is an important step in better understanding how people evaluate others.

These results suggest that diverse contexts may be able to provide quick and powerful tools for prejudice reduction. The studies took participants less than 10 min to complete, and many targets were seen only a few times in flanked arrays before appearing as unflanked stimuli. Thus, it does not appear that the change in attitudes toward the targets is particularly difficult or time consuming to elicit. Indeed, because the decrease in bias is related to automatic associative processes but not to more controlled inhibition processes, situations that reduce a person's motivation or ability to engage in controlled processes, such as a high cognitive load, should not undermine this type of prejudice reduction. This implies that this form of bias reduction may be a viable way to effectively reduce prejudice in people's everyday lives, when they are often operating under a high cognitive load or have little time or motivation to notice and/or resist subtle bias. This also suggests that diverse contexts may be useful across a wide variety of situations. For example, simply viewing an advertising photo that has members of different ethnic groups may be sufficient to begin to decrease bias.

More on mechanism

The modeling results showed that the contexts directly altered the automatic evaluations of the target faces but did not increase the extent to which people effectively inhibited biased responses. The contexts could have affected target evaluations in a number of ways. The diverse flankers may have diminished the extent to which racial categories were activated by the target faces or led to the activation of less stereotypical subtypes of target racial groups (e.g., Wittenbrink et al., 2001). However, if diverse flankers worked by decreasing category accessibility (either directly or through subtyping), then we should not have found a difference in the carry-over effect for familiar and novel faces in Study 2. The fact that the effect of the flankers did not extend to novel group members indicates that the effect was not due to overall decreased category activation, but was, in fact, limited to the specific group members that had been presented with the flankers.

The findings seem most consistent with an evaluative conditioning account. Previous research has found that evaluative conditioning can occur due to a misattribution of affect from one stimulus to another, causing a neutral stimulus to take on the valence of another stimulus presented simultaneously (e.g., Jones et al., 2009; Walther, 2002). Though much of this research has examined the tendency for neutral stimuli to take on the valence of surrounding valenced stimuli, evaluative conditioning can occur when both stimuli are valenced, as well (see Hofmann, De Houwer, Perugini, Baeyens, & Crombez, 2010 for a review). Insofar as racially diverse contexts create an opportunity for evaluative transfer from the context to the focal face, the tendency to implicitly evaluate black faces less favorably than white faces should be decreased when black targets are surrounded by white faces and white targets are surrounded by black faces.

Implications and future directions

Even as these two studies begin to illuminate the effect of social contexts on implicit bias, they suggest a number of important future directions. Interestingly, the results suggest that interpersonal contexts may affect implicit evaluative biases and categorization processes differently. As mentioned in the introduction, research has found that diverse interpersonal contexts lead to an increase in category accessibility and stereotyping (e.g., McGuire et al., 1978; Oakes & Turner, 1986; Stroessner, 1996; Zarate & Smith, 1990). In contrast, the present research found that such contexts have the opposite effect on implicit prejudice.

Future work should investigate more systematically the reasons why cognitive and affective processes appear to be influenced differently by interpersonal contexts, as well as how increased category activation may combine with decreased evaluative bias to affect person perception and interpersonal interactions.

In addition, future research should begin to investigate how more complex information, such as the social relationships among targets in a visual scene, is used in making judgments about those targets. For example, in these studies, participants were given no information about why the targets were presented next to each other. In many situations, though, people are able to judge why individuals are near each other—they may be total strangers, but they may also be members of the same group, or even members from opposing groups (e.g., members of two opposing sports teams playing a game). In situations in which people have such relational information, how does it influence how racially diverse versus homogenous contexts affect prejudice?

Conclusion

Social psychologists have elucidated many of the factors that affect implicit bias toward individual targets. However, research investigating how the presence of other people in the visual field can affect bias has been surprisingly absent. The present research provides initial evidence that racially homogenous contexts increase bias, whereas racially diverse contexts decrease bias, and that these changes in bias are related to changes in automatic evaluations. This research may help to illuminate how bias operates in more complex social situations.

Acknowledgment

This research was supported by a National Science Foundation grant (BCS 0820855) to the second author.

References

- Allen, T. J., Sherman, J. W., & Klauer, K. C. (2010). Social context and the self-regulation of implicit bias. *Group Processes and Intergroup Relations*, 13, 137–150.
- Barden, J., Maddux, W. W., Petty, R. E., & Brewer, M. B. (2004). Contextual moderation of racial bias: The impact of social roles on controlled and automatically activated attitudes. *Journal of Personality and Social Psychology*, 87, 5–22.
- Cohen, B. H. (2001). *Explaining psychological statistics* (2nd ed.). New York: John Wiley and Sons, Inc.
- Conrey, F. R., Sherman, J. W., Gawronski, B., Hugenberg, K., & Groom, C. L. (2005). Separating multiple processes in implicit social cognition: The quad model of implicit task performance. *Journal of Personality and Social Psychology*, 89, 469–487.
- Dasgupta, N., & Greenwald, A. G. (2001). On the malleability of automatic attitudes: Combating automatic prejudice with images of admired and disliked individuals. *Journal of Personality and Social Psychology*, 81, 800–814.
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*, 16, 143–149.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. K. L. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*, 74, 1464–1480.
- Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the implicit association test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology*, 85, 197–216.
- Herr, P. M. (1986). Consequences of priming: Judgment and behavior. *Journal of Personality and Social Psychology*, 51, 1106–1115.
- Herr, P. M., Sherman, S. J., & Fazio, R. H. (1983). On the consequences of priming: Assimilation and contrast effects. *Journal of Experimental Social Psychology*, 19, 323–340.
- Hofmann, W., De Houwer, J., Perugini, M., Baeyens, F., & Crombez, G. (2010). Evaluative conditioning in humans: A meta-analysis. *Psychological Bulletin*, 136, 390–421.
- Jones, C. R., Fazio, R. H., & Olson, M. A. (2009). Implicit misattribution as a mechanism underlying evaluative conditioning. *Journal of Personality and Social Psychology*, 96, 933–948.
- McGuire, W. J., McGuire, C. V., Child, P., & Fujioka, T. (1978). Salience of ethnicity in the spontaneous self-concept as a function of one's ethnic distinctiveness in the social environment. *Journal of Personality and Social Psychology*, 36, 511–520.
- Monteith, M. J., Ashburn-Nardo, L., Voils, C. I., & Czopp, A. M. (2002). Putting the brakes on prejudice: On the development and operation of cues for control. *Journal of Personality and Social Psychology*, 83, 1029–1050.

- Oakes, P., & Turner, J. C. (1986). Distinctiveness and the salience of social category memberships: Is there an automatic perceptual bias towards novelty? *European Journal of Social Psychology*, 16, 325–344.
- Sherman, J. W., Gawronski, B., Gonsalkorale, K., Hugenberg, K., Allen, T. J., & Groom, C. J. (2008). The self-regulation of automatic associations and behavioral impulses. *Psychological Review*, 115, 314–335.
- Stroessner, S. J. (1996). Social categorization by race or sex: Effects of perceived non-normalcy on response times. *Social Cognition*, 14, 247–276.
- Walther, E. (2002). Guilty by mere association: Evaluative conditioning and the spreading attitude effect. *Journal of Personality and Social Psychology*, 82, 919–934.
- Wittenbrink, B., Judd, C. M., & Park, B. (2001). Spontaneous prejudice in context: Variability in automatically activated attitudes. *Journal of Personality and Social Psychology*, 81, 815–827.
- Zarate, M. A., & Smith, E. R. (1990). Person categorization and stereotyping. *Social cognition*, 8, 161–185.