Neural correlates of the motivation to be moral

Nunspeet, F. van

Chapter 5

Controlling implicit prejudice: The effects of moral implications, and evaluation by (non)significant others

This chapter is based on: Van Nunspeet, F., Ellemers, N., & Derks, B. Reducing implicit prejudice against Muslim women: The effects of moral concerns, intra- and intergroup motives. *Manuscript under review.*

The study of attitudes, stereotypes and prejudice is often complicated by social desirability issues: People sometimes adjust their explicit attitudes to appear unbiased (e.g., Crosby, Bromley, & Saxe, 1980). The development of implicit measures of prejudice that capture more automatic biases against social (out)groups was seen to offer a solution to this problem. People may display implicit biases even while they explicitly endorse egalitarian views (e.g., Dovidio, Kawakami, & Beach, 2001), and this is why it is often suggested that implicit prejudice captures the 'automatic' evaluative associations with other groups.

A popular and widely used implicit measure of prejudice is the Implicit Association Test (Greenwald, McGhee, & Schwartz, 1998). The IAT is based on the idea that it is easier to associate your ingroup with positive attributes and an outgroup with negative attributes than vice versa. As a result, participants tend to respond faster on trials in which pictures of ingroup members are associated with positive stimuli (using the same response key) and outgroup members with negative stimuli (congruency). By comparison, they respond more slowly on trials in which ingroup members are associated with negative stimuli and outgroup members with positive stimuli (incongruency). The difference between response latencies on incongruent and congruent trials is taken to assess the degree of implicit bias against a social outgroup.

Although the IAT is frequently presented as a measure of automatic bias, by now several studies have shown the malleability of 'automatic prejudice. This suggests that implicit biases can be influenced too, for example by self-concerns and social motives (for an overview see Blair, 2002). Effects of self-concerns are shown in research were the induction of stereotype threat among Whites –by triggering the stereotype that they are racists– increased implicit biases towards Blacks (Frantz, Cuddy, Burnett, Ray, & Hart, 2004; Rudman, Dohn, & Fairchild, 2007). Other research has revealed that implicit biases can also be affected by intergroup concerns: When a Black experimenter was present during participants' performance on an IAT, Whites were able to inhibit their pro-White bias (e.g., Lowery, Hardin, & Sinclair, 2001). Additionally, Richeson and Ambady (2003) showed the significant effect of the role of such a Black person present: Their participants also displayed a smaller bias against Blacks, but only when their Black

partner in the experiment was a superior (instead of a subordinate). Furthermore, Van Nunspeet, Ellemers, Derks, and Nieuwenhuis (2014) revealed that emphasizing the moral implications of performance on an IAT –compared to implications concerning individual competence– led participants to show a smaller bias against Muslims. In a follow-up study, this effect was particularly strong when people were evaluated by minimal ingroup (rather than outgroup) members, thus suggesting effects of intragroup concerns (Van Nunspeet, Derks, Ellemers, & Nieuwenhuis, under review).

Current research

Although different motives and contexts have been shown to affect people's evaluative bias, to our knowledge they have not been directly compared in one study. It is thus unclear which concern or motive would benefit the control of bias against an outgroup when for example, interpersonal contact with a person from the target group is not feasible. In the current research, our aim is to examine the effects of three different interventions on people's ability to control their evaluative bias against an outgroup in one IAT experiment: (1) Personal concerns about moral implications of displaying bias; (2) intergroup motives (i.e., concerns about displaying bias in front of a representative of the devalued group) and (3) intragroup concerns about displaying bias in front of self-relevant others. Specifically, we demonstrate how people's evaluative bias against Muslims is affected by (1) emphasizing the moral (compared to competence) test implications of the IAT; (2) having participants be observed by either a Muslim or a non-Muslim evaluator (first ingroup/outgroup dimension); and (3) presenting this evaluator as either a minimal ingroup or outgroup member (second ingroup/outgroup dimension, resulting in cross-categorization). In the current study we combined these interventions to directly compare their effects on reducing implicit evaluative bias and to examine whether and how they may influence one another.

Additionally, we aimed to examine the underlying processes associated with reducing implicit bias. In studies concerning the effects of personal and social motives on people's evaluative biases, little attention has been devoted to how such a bias (i.e., IAT performance) was affected. In an IAT, bias is reduced by

diminishing the difference between response latencies on stereotype-incongruent and stereotype-congruent trials. However, this can be accomplished in two ways: Either by becoming quicker on incongruent trials (and thus becoming better in associating the outgroup with positive attributes), or by responding more slowly on congruent trials (and inhibiting negative associations with the outgroup and positive associations with the ingroup). Interestingly, the smaller IAT effect in research of Richeson and Ambady (2003) was due to slower responses on congruent trials. In a similar vein, Van Nunspeet et al. (2014) revealed that an emphasis on morality caused participants to show a smaller IAT bias, caused by their slowed down responses on congruent trials. In addition, these researchers showed that stressing the moral test implications was associated with enhanced response-monitoring (measured using EEG). results suggested that participants' reduced bias was related to the inhibition of prepotent responses on stereotype-consistent (i.e., congruent) trials (Van Nunspeet et al., 2014). In the current research, we therefore examined the pattern of response latencies on congruent and incongruent trials separately to see how exactly the three types of interventions affected participants' evaluative bias.

Study 5.1

Method

Participants.

Only female, non-Muslim, students (N = 225; M_{age} = 20.5 years, SD = 2.6) participated in the study and received either money or course credit for their participation. Two participants were excluded from analyses: One due to technical problems, another because she responded too late on all IAT trials, indicating lack of attention. Participants were randomly assigned to one of the eight experimental conditions of the 2 (Task Domain: morality/competence) x 2 (Evaluator's Minimal Group: ingroup/outgroup) x 2 (Evaluator's Religion: Muslim/non-Muslim) between-participants design. Note that the evaluator was the same individual in all conditions, but that she did or did not wear a headscarf (see Figure 5.1).

Procedure.

Participants were seated in an individual computer room with a webcam on top of the computer screen, and a camera behind them in a top corner of the

cubicle. They were told that they would be working together with another participant. They then completed a (bogus) questionnaire that was said to assess whether they had either a so-called 'P'- or 'O'- personality style. After a short waiting period, participants learned about their own alleged personality style and the styles of the other participants and they were informed whom they would be working with during the experiment. The other person either was said to have the same personality style as the participant (to convey this individual was a member of the same minimal group as the participant), or she allegedly had the other personality style (to indicate this individual belonged to a different group). Participants then read that they would perform a computer task. During the first part of the experiment, the other person would supposedly observe and give them feedback after every trial and the roles would be reversed in the second part. Thereafter, a webcam connection was simulated: The other person introduced herself and said that she would observe and provide visual feedback on every trial. Then, participants read either the morality or competence instruction and started with the IAT. In reality, all participants were said to have a 'P'- personality style and were introduced to a confederate whose movies were prerecorded. Feedback displays during the IAT were related to participants' actual responses (i.e., positive feedback when they responded correctly, negative feedback when they responded incorrectly). After the IAT, participants completed some self-report items and were properly debriefed.

Task domain manipulation. Before the start of the IAT, half of the participants read that the computer task they were going to perform could indicate their endorsement of *moral values* concerning egalitarianism and discrimination (the morality condition). The other half of the participants was informed that the test could indicate their *ability* to process new information and to learn new tasks (the competence condition). All participants were instructed to respond as quickly and accurately as possible and the test implications were repeated before the start of each test block (see also Van Nunspeet et al., *2014*).

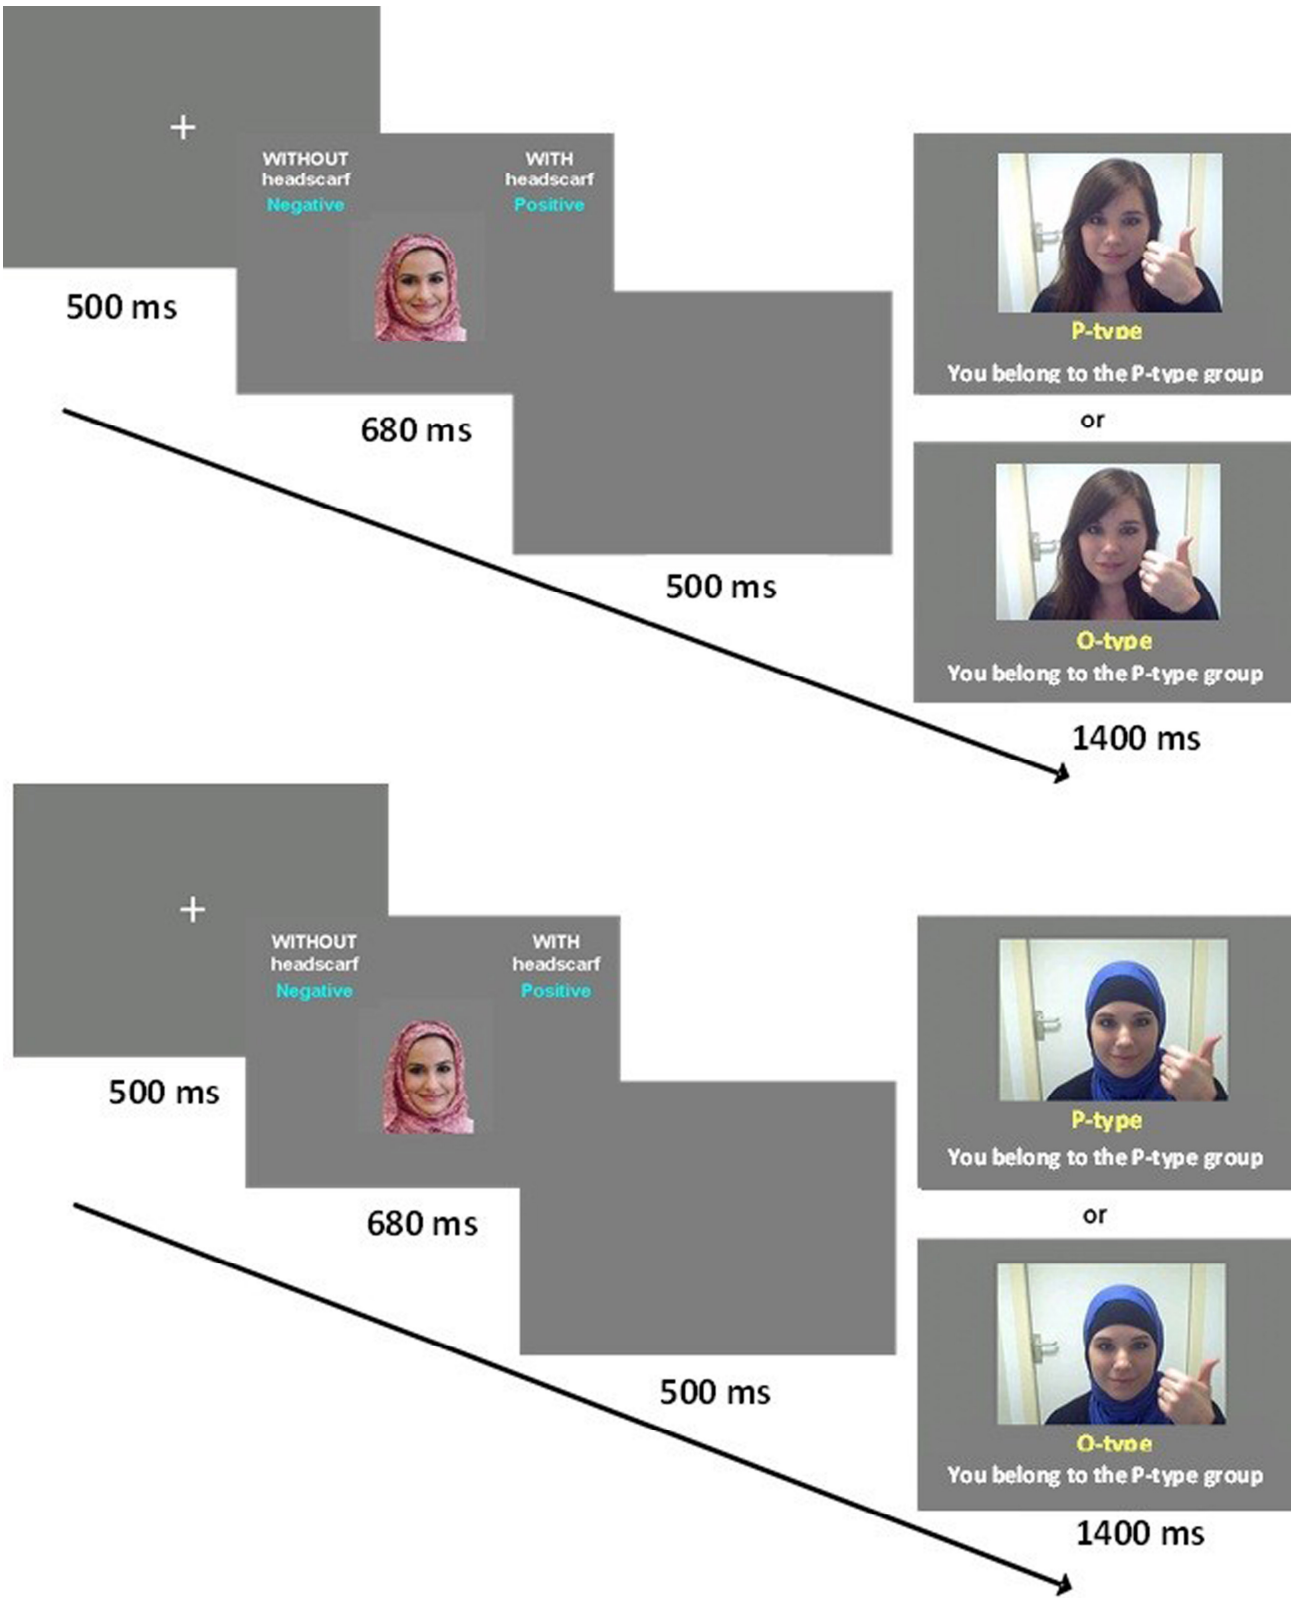

Figure 5.1. Example of an (incongruent) IAT trial. The (same) evaluator resembled either a non-Muslim (top) or Muslim (bottom) woman.

Instruments.

The Implicit Association Test. Participants performed the five blocks of the IAT as designed by Greenwald et al. (1998). Stimuli representing the target concepts consisted of 10 pictures of Muslim women (wearing a headscarf) and 10 pictures of non-Muslim women (not wearing a headscarf). Stimuli that represented positive and negative attributes consisted of 5 pictures of positive scenes, and 5 pictures of negative scenes, selected from the International Affective Picture System (Lang et al., 2005).

In (training) block 1, participants were asked to respond to the pictures of women by pressing a left key for Muslim women and a right key for non-Muslim women. In (training) block 2 they were asked to use the same two keys to respond to the negative and positive pictures. In block 3 (a test block) both picture types were presented and participants responded with one key to pictures of both Muslim women and negative scenes and with the other key to pictures of both non-Muslim women and positive scenes (i.e., congruent trials). In (training) block 4, the response keys for the pictures of (non-)Muslim women were switched and in block 5 (a test block), participants had to respond to pictures of both non-Muslim women and negative scenes with one key and to pictures of both Muslim-women and positive scenes with one other key (i.e., incongruent trials). Blocks 1, 2 and 4 consisted of 20 trials, blocks 3 and 5 of 70 trials each. Every trial started with a fixation point (500 ms), followed by stimulus presentation (680 ms), a blank screen (500 ms) and a feedback screen (1400 ms). The feedback screen consisted of a movie clip of the evaluator showing either positive (smiling and holding 'thumbs up') or negative (frowning and pointing 'thumbs down') feedback. To ensure that participants were aware of the minimal group membership of their evaluator, we inserted a text display below the movie indicating the personality type of the evaluator, and a text display at the bottom of the screen indicating the personality type group of the participant (see Figure 5.1). In case participants did not respond in time, they saw the words "too late".

The IAT effect. The dependent measure was the IAT effect, indicated by the *D* score, and measured as the difference in reaction times on incongruent and congruent trials divided by a pooled *SD* of all correct trials (according to the

scoring algorithm described by Greenwald et al. 2003). We included all trials, replaced error latencies with a replacement value (M + 2 $SD_{correct}$) and replaced latencies exceeding the maximum response time with the maximum response time of 680 ms. The resulting positive D scores are an indication of people's evaluative bias against Muslim women.

Checks. Directly after the IAT, we checked the task domain manipulation: Participants were asked to indicate what the IAT intended to measure. They could indicate that the test either measured how well they were able to process information and to learn new tasks, or that it assessed their moral values concerning egalitarianism and discrimination. Second, we checked the evaluator's minimal group manipulation by asking participants to indicate whether their evaluator was a member of the same or another minimal group. Furthermore, we tested participants' perceptions of the validity of the test (i.e., "My test score can assess what kind of person I am"), and their overall impression of their evaluator ("I think the participant who gave me feedback is competent/kind/moral", 3 items). Participants could respond on a 7-point Likert scale (1 = completely disagree, 7 = completely agree).

Results

Checks.

Results concerning the manipulation of task domain showed that 96% (N = 105) of participants in the morality condition indicated that the test measured their moral values concerning egalitarianism and discrimination. Moreover, ninety-seven percent (N = 110) of participants in the competence condition indicated that the test measured their ability to quickly process information and learn new tasks. Results concerning the evaluator's minimal group manipulation showed that 95% (N = 103) of participants whose evaluator was an ingroup member correctly answered that their evaluator was a member of their own group. One hundred percent (N = 115) of participants whose evaluator was an outgroup member answered correctly that their evaluator was a member or the other group. Excluding the participants who answered one of the checks incorrectly (N = 10) did not alter the pattern of the means. We therefore included those participants in all analyses.

The perceived validity of the IAT and participants' impression of their evaluator showed that, as intended, there were no reliable effects of experimental condition on participants' perceived validity of the test (overall M = 3.32, SD = 1.48; F's ≤ 2.71, p's ≥ .10) or their impression of their evaluator, which was quite positive overall ($M_{competent}$ = 5.17, SD = 1.14; M_{kind} = 5.70, SD = 0.87; M_{moral} = 5.24, SD = 0.97; all F's ≤ 3.87, p's ≥ .06).

IAT effect (*D* score).

An ANOVA with task domain, evaluator's minimal group and evaluator's religion as independent factors revealed a significant main effect of evaluator's religion, $F(1,215) = 11.68$, $p = .001$, $\eta_p^2 = .05$. Whereas participants whose evaluator was a non-Muslim woman showed significant bias against Muslim women (M = 0.16, SD = 0.45; $t[108] = 3.73$, $p < .001$), this bias was reduced to non-significance when participants were evaluated by a Muslim woman (M = -0.04, SD = 0.45, $t[113] = -0.90$, $p = .37$). Additionally, the interaction between task domain and evaluator's religion was marginally significant, $F(1,215) = 2.88$, $p = .09$, $\eta_p^2 = .01$. Analysis of simple main effects indicated that when evaluated by a Muslim woman there was no difference in IAT bias between the morality and competence condition (M = -0.03, SD = 0.50, M = -0.04, SD = 0.41 respectively; $F < 1$). However, when evaluated by a non-Muslim woman, participants for whom the moral implications of the test were emphasized showed a significantly weaker negative bias (M = 0.07, SD = 0.46) than participants for whom the implications of the test concerning their competence were emphasized (M = 0.27, SD = 0.42), $F(1,215) = 4.99$, $p = .03$, $\eta_p^2 = .02$. These results show that having a Muslim evaluator present is an impactful way of reducing non-Muslims' implicit anti-Muslim bias. However, even in the absence of an evaluator from the target group, a focus on morality rather than competence also reduces implicit bias significantly[10].

[10] A prior study (Van Nunspeet et al., under review) showed that emphasizing morality rather than competence reduced implicit bias in the presence of a (non-Muslim) evaluator belonging to a minimal ingroup, but not when this evaluator belonged to a minimal outgroup. Although this interaction effect was not significant in the current study ($F < 1$), the effect of task domain was indeed stronger when participants thought they were evaluated by a minimal ingroup member ($M_{morality}$ = 0.04, SD = 0.52; $M_{competence}$ = 0.27, SD

Inspection of reaction times.

To examine whether the effects of evaluator's religion and task domain on implicit bias were due to enhanced positive associations with the Muslim outgroup (reduced RTs on incongruent trials) or the inhibition of prepotent biased responses (increased RTs on congruent trials), we analyzed response latencies on correctly answered congruent and incongruent trials separately.

Congruent trials. The analysis of response latencies on correct *congruent* trials (reflecting the speed of making stereotype-congruent associations) revealed significant effects of our manipulations in line with the observed pattern of implicit bias reduction reported above. Parallel to the effect of evaluator's religion on the implicit bias score, evaluator's religion significant affected RTs on congruent trials, $F(1,215) = 7.09$, $p = .008$ $\eta_p^2 = .03$. Participants whose evaluator was a Muslim woman responded more slowly on congruent trials ($M = 503.97$, $SD = 24.24$) than participants whose evaluator was a non-Muslim woman ($M = 495.45$, $SD = 27.13$). Moreover, replicating previous work (Van Nunspeet et al., 2014), participants working under moral task instructions responded significantly more slowly on congruent trials ($M = 502.81$, $SD = 24.33$) than participants in the competence condition;($M = 496.88$, $SD = 27.30$), $F(1,215) = 3.92$, $p = .05$, $\eta_p^2 = .02$. Finally, participants responded marginally slower on congruent trials when their evaluator was a minimal ingroup member ($M = 502.89$, $SD = 26.64$) than when she was a minimal outgroup member ($M = 496.91$, $SD = 25.14$), $F(1,215) = 2.73$, $p = .10$, $\eta_p^2 = .01$.

Although there were no significant interaction effects; F's ≤ 1.84, $p \geq .18$, to enable a more direct comparison with the analyses for overall implicit bias, we analyzed RTs on congruent trials per evaluator's religion condition. Replicating the pattern for implicit bias, when participants were evaluated by a Muslim woman there were no significant effects of task domain or evaluator's minimal group on congruent response latencies (F's ≤ 2.44, p's $\geq .12$). However, when evaluated by a non-Muslim woman, participants responded significantly slower on congruent trials

---

$= 0.52$; $F[1,105] = 3.73$, $p = .06$, $\eta_p^2 = .03$), compared to a minimal outgroup member ($M_{morality} = 0.11$, $SD = 0.40$; $M_{competence} = 0.26$, $SD = 0.31$, $F[1,105] = 1.49$, $p = .23$).

in the morality condition (M = 500.61, SD = 23.66) than in the competence condition (M = 489.36, SD = 29.83), $F(1,105) = 4.67$, $p = .03$, $\eta_p^2 = .04$.

Incongruent trials. Analysis of response latencies on the correct incongruent trials (reflecting the stereotype-incongruent combinations of Muslims/positive and non-Muslim/negative) revealed no main effects of task domain, evaluator's religion or evaluator's group type, nor the interaction between evaluator's religion and task domain found for the overall D-score (all F's $\leq$ 1.04, $p \geq .31$). Thus, the experimental manipulations that resulted in a reduction of implicit bias did not cause participants to respond faster on incongruent trials[11].

Discussion

The results of Study 5.1 showed that participants reduced their anti-Muslim bias in case of presence of a Muslim evaluator or, in the absence of a Muslim evaluator, the emphasis on their morality instead of their competence. Moreover, this bias reduction was associated with the inhibition of stereotype conforming responses rather than with increased positive associations with the Muslim outgroup. Although these findings are consistent with previous research (Richeson & Ambady, 2003; Van Nunspeet et al., 2014), we wanted to test whether they are dependent upon the duration of the experiment: If positive associations have to be learned, they may only develop over a longer period of time.

We examined this possibility in Study 5.2, in which we increased the exposure to participants' evaluator while using the same cross-categorization dimensions as in Study 5.1. If participants share their minimal group membership with their Muslim evaluator, they may become to perceive their evaluator as a partial ingroup member when the duration of the interaction is increased (see also Crisp & Hewstone, 1999; Crisp, Hewstone, & Rubin, 2001, for effects of cross-categorization). Moreover, perceiving the evaluator as a partial ingroup member

[11] We also found an unexpected interaction between task domain and evaluator's group type, $F(1,215) = 4.02$, $p = .05$, $\eta_p^2 = .02$. Whereas there was no difference between the minimal group types of the evaluator in the morality condition ($M_{ingroup}$ = 495.28, SD = 24.98; $M_{outgroup}$ = 498.33, SD = 22.66, $F < 1$), participants in the competence condition responded faster on incongruent trials when the evaluator was a minimal outgroup (M = 490.10, SD = 21.74) instead of a minimal ingroup member (M = 499.61, SD = 22.78), $F(1,215) = 4.63$, $p = .03$, $\eta_p^2 = .02$.

may facilitate positive associations with the Muslim outgroup. In Study 5.2, we thus significantly increased the number of IAT trials to enable participants to develop new (positive) associations with Muslims during the task (resulting in reduced RTs on incongruent trials).

Study 5.2

Method

Participants.

Only female, non-Muslim, students (N = 102; M_{age} = 21.3 years, SD = 3.1) participated in the study for money or course credits. One participant was excluded from the analyses because she responded too late on more than 25% of the IAT trials, suggesting lack of attention to the experimental task.

Procedure.

The IAT and the procedure were similar to those described in Study 5.1. However, in Study 5.2, all participants received feedback from a Muslim evaluator. Thus, participants were randomly assigned to one of the four experimental conditions of the 2 (Task Domain: morality/competence) x 2 (Evaluator's Minimal Group: ingroup/outgroup) between-participants design. Moreover, the amount of trials in the two test blocks of the IAT was increased: From 70 trials per block in the previous study to 120 trials per block in Study 5.2.

Results

Checks.

Ninety-eight percent (N = 49) of participants in the morality condition and 96% (N = 49) of participants in the competence condition correctly reported the task domain. Moreover, 92% (N = 47) of participants whose evaluator was an ingroup member and 98% (N = 49) of participants whose evaluator was an outgroup member reported their evaluators' minimal group correctly. Because exclusion of the participants who answered one of the checks incorrectly (N = 6) did not alter the pattern of means, we included those participants in all analyses.

As intended, participants in all four conditions indicated that the test was able to assess what kind of person they are to a similar degree; overall M = 3.44, SD = 1.57; F's ≤ 1.23, p's ≥ .27. Moreover, there were no effects of our task domain or evaluator's minimal group manipulation on participants' impression of their

evaluator, which was quite positive overall ($M_{competent}$ = 5.36, SD = 1.09; M_{kind} = 5.85, SD = 0.84; M_{moral} = 5.54, SD = 0.98; all F's ≤ 1.68, p's ≥ .20).

IAT effect (*D* score).

Consistent with Study 5.1, now that all participants were evaluated by a Muslim woman, on average they did not show implicit bias against Muslim women, M = -.02, SD = .32, $t(100)$ = -.53, p = .60. Additionally, an ANOVA with task domain and evaluator's minimal group type as independent factors revealed a main effect of evaluator's minimal group type: Participants whose Muslim evaluator was presented as a minimal ingroup member showed significantly less bias against Muslim women (M = -0.08, SD = 0.27) compared to participants who thought they were evaluated by an outgroup member (M = 0.05, SD = 0.35), $F(1,97)$ = 5.02, p = .03, η_p^2 = .05. The effect of task domain was marginally significant, $F(1,97)$ = 2.89, p = .09, η_p^2 = .03: In line with the previous findings the means show that implicit bias was reduced under moral task instructions (M = -0.07, SD = 0.33) compared to competence instructions (M = 0.03, SD = 0.29).

We proceeded by examining whether RTs on correct congruent and incongruent trials differed across experimental conditions. Interestingly, the general tendency to slow down on congruent trials indicating the inclination to inhibit prejudice conforming responses did not depend on the evaluator being an in- or an outgroup member or on task domain (F's ≤ 2.66, p's ≥ .11). Additionally, and as expected, we found evidence in line with our reasoning that increasing the number of trials in which participants are exposed to a Muslim evaluator who is presented as an ingroup member can facilitate the ability to associate positive stimuli with Muslim targets. That is, participants responded faster on incongruent trials when the Muslim evaluator was presented as a minimal ingroup member (M = 478.87, SD = 23.90) than when she was an outgroup member (M = 493.26, SD = 23.08), $F(1,97)$ = 9.47, p = .003, η_p^2 = .09[12]. This suggests that the decrease in implicit bias

[12] To directly test the effect of the increase in trials, we combined the data of Study 5.2 (N = 101) with the data of participants who were evaluated by a Muslim evaluator in Study 5.1 (N = 114). Results of an ANOVA with RTs on incongruent trials as dependent variable and amount of trials, task domain and evaluator's minimal group type as independent factors showed a main effect of amount of trials: Participants responded significantly faster on incongruent trials when the amount was increased

observed when the Muslim evaluator was a minimal ingroup member reflects that the ability to associate Muslim individuals with positive stimuli is facilitated under these conditions.

General Discussion

In the current research we directly compared the effects of three different interventions on people's implicit evaluative bias against Muslims: (1) People's personal motives to appear moral; (2) their intergroup motivation to perform well towards a Muslim evaluator, and (3) their intragroup-based motives to perform well in front of self-relevant others (categorized on a second, minimal, group dimension). We tested these effects by introducing a Muslim/non-Muslim IAT as a measure of participants' moral values or of their competence. Moreover, participants performance was evaluated by either a non-Muslim or Muslim individual who was presented as a minimal in- or outgroup member. Results of Study 5.1 revealed the significant effect of target presence: In line with previous research (Lowery et al., 2001), participants showed no sign of anti-Muslim bias when they their evaluator was Muslim. Moreover, the significant reduction in bias was associated with the inhibition of prejudice: Instead of decreased response times on incongruent trials (indicating rapid associations between Muslims and positive attributes and non-Muslims and negative attributes), participants slowed down their responses on congruent trials, suggesting that they aimed to inhibit their prepotent responses to rapidly associate Muslims with negativity and non-Muslims with positivity.

In case participants' evaluator was not Muslim, we did find the same pattern of inhibition of prejudice-conforming responses when the moral implications of the test were emphasized: When participants were told that their test score could be perceived as an indication of their moral values concerning egalitarianism, this helped them to show a smaller bias against Muslims than when they were told that

---

$M_{120trials}$ = 485.99, SD = 24.47, $M_{70trials}$ = 496.90, SD = 22.66, F(1,207) = 11.82, p = .001, η_p^2 = .05. Moreover, there was a significant interaction effect between amount of trials and evaluator's minimal group type (F[1,207] = 7.50, p = .007, η_p^2 = .04), indicating that participants only responded faster on incongruent trials while they were evaluated by a minimal ingroup member in case of the increased amount of trials; $M_{120trials}$ = 478.86, SD = 23.90, $M_{70trials}$ = 498.75, SD = 22.64, F(1,207) = 18.82, p < .001, η_p^2 = .08.

their test could reveal their competence. Emphasizing one' morality thus seems to be an effective way to facilitate bias reduction and may be an alternative intervention when intergroup contact is not feasible.

Furthermore, in line with previous research (e.g., Kawakami, Dovidio, Moll, Hermsen, & Russin, 2000), results of Study 5.2 revealed that new (positive) associations can also be induced. First of all by increasing the amount of exposure to a Muslim evaluator and thus by emphasizing that one's intergroup behavior is evaluated by an outgroup member. And second, by introducing cross-categorization and focusing people on what they have in common with someone who they perceive as an outgroup member on another social dimension: Presenting a Muslim (outgroup) evaluator as a minimal ingroup member helped participants to developed positive associations with the Muslim outgroup. Importantly, our results extend prior research which revealed that shared (minimal) group membership(s) can override people's explicit evaluative bias against outgroup members (e.g., Crisp et al., 2001; Urada, Stenstrom, & Miller, 2007), by showing similar findings for people's implicit bias.

Our findings indicate that there are different ways in which implicit prejudice can be reduced. The presence of a member of the target outgroup may have the greatest impact on the control of prejudiced responses and can even activate new (positive) associations with the outgroup. However, we should not overestimate this effect in everyday interactions: Social groups that are the focus of prejudice research are generally minority groups in society that are often segregated from the majority in education, housing, and work, preventing extensive intergroup interactions. The current research thus offers a contribution to insights on prejudice reduction by demonstrating again the potential impact of emphasizing one's morality and the presence of others who share the same ingroup norms, even when no outgroup member is present (see also Van Nunspeet et al., 2014).

We note that specific circumstances were in place in the current research as it remains unclear which aspect of our manipulations concerning the Muslim evaluator caused the effect of faster positive associations with Muslim women. Our participants received feedback on every trial and since they made few errors, they received almost continuous positive feedback. They thus repeatedly saw a smiling,

approving Muslim woman who was presented as someone like them (an ingroup member). It is less likely that similar effects will be obtained when participants were provided with as much or more negative rather than positive feedback.

Nevertheless, we have shown that evaluative bias against Muslims can be reduced by several means. Presence of a Muslim evaluator causes people to inhibit their prejudiced responses and, provided there is enough exposure, presenting her as a self-relevant other may strengthen positive associations. Moreover, besides this form of intergroup contact, prejudice control can also be instigated by emphasizing people's moral values.

Acknowledgements

We thank Johannes Parzonka, Eva Schildkamp, and Titia van Malestein for help with the data collection; Bianca van Nunspeet for her contribution to the experiment; and Kees Verduin for his help with processing the data.