Journal of Personality and Social Psychology

Biased Allocation of Faces to Social Categories

Ron Dotsch, Daniël H. J. Wigboldus, and Ad van Knippenberg

Online First Publication, March 28, 2011. doi: 10.1037/a0023026

Journal of Personality and Social Psychology

0022-3514/11/$12.00 DOI: 10.1037/a0023026

Biased Allocation of Faces to Social Categories

Ron Dotsch, Daniël H. J. Wigboldus, and Ad van Knippenberg
Radboud University Nijmegen

Three studies show that social categorization is biased at the level of category allocation. In all studies, participants categorized faces. In Studies 1 and 2, participants overallocated faces with criminal features—a stereotypical negative trait—to the stigmatized Moroccan category, especially if they were prejudiced. On the contrary, the stereotype-irrelevant negative trait stupid did not lead to overallocation to the Moroccan category. In Study 3, using the stigmatized category homosexual, the previously used negative trait criminal—irrelevant to the homosexual stereotype—did not lead to overallocation, but the stereotype-relevant positive trait femininity did. These results demonstrate that normative fit is higher for faces with stereotype-relevant features regardless of valence. Moreover, individual differences in implicit prejudice predicted the extent to which stereotype-relevant traits elicited overallocation: Whereas more negatively prejudiced people showed greater overallocation of faces associated with negative stereotype-relevant traits, they showed less overallocation of faces associated with positive stereotype-relevant traits. These results support our normative fit hypothesis: In general, normative fit is better for faces with stereotypical features. Moreover, normative fit is enhanced for prejudiced individuals when these features are evaluatively congruent. Social categorization thus may be biased in itself.

Keywords: social categorization, prejudice, face processing, implicit association, stereotype

Supplemental materials: http://dx.doi.org/10.1037/a0023026.supp

As much as we would like to individuate everyone we meet, we effortlessly and automatically categorize persons into groups to simplify and make sense of the enormous amount of social information in the world (Allport, 1954). Although much research exists on the consequences of categorization (e.g., Fiske, 1998; Tajfel, 1969; Turner, 1987), our knowledge about the category selection process itself is rather limited. We cannot claim full understanding of the categorization phenomenon "unless we also know who gets placed into which categories and why" (Zebrowitz, 1996, p. 80). Processing goals, motivational states, and (chronic) accessibility have been identified as some of the determinants of category selection and paint a clear picture of why specific categories have a higher chance of being selected when a target person belongs to multiple categories (Hugenberg & Sacco, 2008; Macrae, Bodenhausen, Milne, & Calvini, 1999; Macrae, Bodenhausen, Milne, Thorn, & Castelli, 1997; van Knippenberg, van Twuyver, & Pepels, 1994). However, before this selection can be made, the cognitive system has to determine which categories pertain to a perceived person in the first place. In the current article we test the prediction that at this early stage of social categorization, the very process of category allocation is biased.

Ron Dotsch, Daniël H. J. Wigboldus, and Ad van Knippenberg, Behavioural Science Institute, Radboud University Nijmegen, the Netherlands.

We are grateful to Gijs Bijlstra for gathering the pilot data about the Dutch Moroccan stereotype and to Giti Bakhtiari for assistance with data collection.

Correspondence concerning this article should be addressed to Ron Dotsch, who is now at Department of Psychology, Princeton University, Green Hall, Princeton, NJ 08544-1010. E-mail: rdotsch@princeton.edu

Category allocation was investigated as early as the late 1950s. Pettigrew, Allport, and Barnett (1958) conducted a binocular rivalry experiment in South Africa in which they briefly presented, among other combinations, a White face to one eye and a face with darker skin tone (colored, African, or Indian) to the other eye. Despite the dual stimulus, participants perceive this as one single face. Afrikaners, that is, White people who in the past were generally in favor of racial separation (Dubow, 1992), were more prone to allocate any combination of White and non-White stimulus faces to a non-White category than English White, African, or Indian participants. Pettigrew et al. explained these findings as an effect of prejudice on the allocation of ambiguous exemplars to racial categories. These data constitute preliminary evidence for biases in category allocation.

Other attempts at revealing biases in category allocation have focused solely on accuracy, for example, by investigating the influence of anti-Semitism on the categorization of Jewish and non-Jewish faces (e.g., Allport & Kramer, 1946; Elliott & Wittenberg, 1955; Himmelfarb, 1966; Quanty, Keats, & Harkins, 1975), the categorization of multiple ethnicities in a multicultural context (Kosic & Phalet, 2006), and the effect of ingroup identification on categorization (Castano, Yzerbyt, Bourguignon, & Seron, 2002). These studies showed that a bias existed for more prejudiced participants to allocate more faces to a stigmatized category than less prejudiced participants, although no differences in accuracy were observed. However, in our view, this focus on accuracy may have obscured more fine-grained category allocation biases. Although extant research has shown that more prejudiced people allocate more faces to a stigmatized category than do less prejudiced people, it remains unclear on the basis of which facial features this larger number of category allocations occurs. As we

argue below, this greater allocation is not random but based on specific facial features that are more diagnostic of stigmatized category membership in the minds of prejudiced people than in the minds of nonprejudiced people. In the present study, we aim to demonstrate these previously unnoticed specific biases in category allocation.

When does a category pertain to a perceived person? It has been suggested that there has to be normative fit between that person and the category (Bruner, 1957; Oakes, 1987; van Knippenberg & Dijksterhuis, 2000). Normative fit refers to the match between stimulus characteristics and the characteristics a perceiver expects members of specific categories to have. For instance, when the features or configuration of a perceived face match someone's expectations of what typical Chinese faces look like, there is normative fit. The stereotype is the cognitive structure containing these expectations (Stangor & Schaller, 1996; van Knippenberg & Dijksterhuis, 2000). Importantly, this means that normative fit is affected by the idiosyncratic content of the stereotype. If a culturally shared stereotype holds that members of a specific category are criminal, a person endorsing this stereotype will expect category members to look criminal, and therefore, normative fit will be enhanced when this person perceives a criminal-looking person.

There can be individual differences in the extent to which people endorse culturally shared stereotypes: Although some might believe strongly that members of a social category are criminal, others might have weaker associations between the category and criminality. Importantly, this variation across individuals is predicted by their level of prejudice, that is, the extent to which they have a negative evaluation of the category as a whole (Gordijn, Koomen, & Stapel, 2001; Wittenbrink, Judd, & Park, 1997). To the extent that highly prejudiced people believe category members to be criminal, they will expect members of that group to look more criminal. This relationship recently has been established within the domain of faces. Dotsch, Wigboldus, Langner, and van Knippenberg (2008) used a forced-choice paradigm to reconstruct what people believed typical Moroccan faces looked like (Moroccans are a highly stigmatized immigrant group in the Netherlands; see Coenders, Lubbers, Scheepers, & Verkuyten, 2008; Verkuyten & Zaremba, 2005; and strongly associated with the trait criminal; Gordijn et al., 2001). Images produced by highly prejudiced people depicted faces that were more criminal-looking than those produced by people low in prejudice. Because highly prejudiced individuals expect Moroccan people to be more criminal, criminal-looking faces have better normative fit for the Moroccan category, and, therefore, highly prejudiced individuals should be more likely to categorize criminal-looking stimulus faces as Moroccan.

Alternatively, Ruys, Dijksterhuis, and Corneille (2008) proposed an evaluative fit hypothesis such that social categorization is facilitated for exemplars that are evaluatively congruent with the target category. With regards to the Moroccan category, this means that because highly prejudiced individuals evaluate the Moroccan category more negatively, negative-looking faces have better evaluative fit for the Moroccan category; therefore, highly prejudiced individuals should be more likely to categorize criminal-looking stimulus faces as Moroccan. The difference between evaluative fit and normative fit is that the first concerns only valence congruity due to prejudice, whereas the latter involves stereotypicality. It is often difficult to disentangle prejudice-related valence from stereotyping (e.g., Amodio & Devine, 2006; Wittenbrink et al., 1997), because for the majority of social categories, the two concepts seem to be confounded: For example, the trait criminal—which is stereotypical for Moroccans—is negative, and the Moroccan category itself is negatively evaluated. However, in our view, prejudice and stereotyping have clearly identifiable contributions to fit.

Evaluative fit cannot possibly account for the vast number of potential social categories that can be allocated given a stimulus of some valence. The unidimensional concept of valence in the majority of cases is not predictive for category membership: Not all negatively evaluated people are members of one specific negative category. Fit by necessity depends additionally on a concept of higher dimensionality: the stereotype. The stereotype contains those traits that are predictive of group membership (see Le Pelley et al., 2010). Some of those traits might have the same valence as the category; others might not. Therefore, in our view, faces with features associated with traits that are predictive of group membership (i.e., stereotypical traits) will have fit, regardless of valence. If a perceiver has a negative evaluation of that category (i.e., is prejudiced), faces with features that are both stereotypical and negative will have enhanced fit.

We propose that faces with negative features will have enhanced fit for prejudiced people only to the extent that these features are associated with stereotypical traits. This hypothesis differs from the evaluative fit hypothesis in that prejudice does not enhance fit for faces with features associated with any negative trait, but does so on the condition that this negatively valenced trait is stereotypical. Importantly, this specific definition of normative fit differs also from the more simplistic definition that prejudice enhances fit for faces with more features associated with stereotypical traits, regardless of the valence of those traits. In three studies, the present article aims to test this normative fit hypothesis. Study 1 aims to show that people indeed overallocate faces with features associated with negative stereotype-relevant traits to a negatively evaluated category and that this happens to a greater extent for more prejudiced people. By the *overallocation* of faces, we mean that participants allocate more faces with a specific set of features to a category relative to faces without those features. We use the term *underallocation* for the opposite effect, that is, for the allocation of fewer faces with a specific set of features to a category relative to faces without those features.[1] Furthermore, in Study 1 we aim to show that only stereotype-relevant negative traits, but not stereotype-irrelevant negative traits, elicit overallocation. In Study 2, we aim to rule out several alternative explanations. Finally, in Study 3 we aim to show that even valence-incongruent traits may elicit overallocation on the condition that these traits are stereotype relevant, but then prejudiced participants will not display enhanced overallocation.

Study 1

Study 1 was designed to show that more prejudiced people allocate more faces with features associated with a negative stereotype-relevant trait (criminal) to a stigmatized category (Moroccans) than do less prejudiced people. Moreover, with respect to stereotype-irrelevant traits, in Study 1 we also aimed to show that

[1] By no means do we want to suggest any accuracy-related connotation in using the terms *underallocation* and *overallocation*, in the sense that they would imply bias with respect to accurate allocations.

more prejudiced people do not allocate more faces with features associated with a negative stereotype-irrelevant trait (in this case, stupid) to the same stigmatized category, thereby demonstrating that normative fit is not enhanced by just any evaluatively congruent trait. To this end, Dutch participants were instructed to categorize faces as Moroccan or non-Moroccan. The faces in the categorization task were taken from Dotsch et al. (2008, Study 2). A number of these faces were manipulated to have either more criminal-looking features or more stupid-looking features than the original faces. After the categorization task, implicit prejudice was measured.

Criminality is a highly stereotype-relevant trait for Moroccans: When asked to list the content of the cultural stereotypes of Moroccans, Dutch participants most frequently listed the trait criminal (70%; Gordijn et al., 2001). The negative trait stupid was mentioned by fewer than 20% of participants and was therefore not considered a part of the shared cultural stereotype of Moroccans. Because normative fit is enhanced for faces with features associated with stereotype-relevant traits, we expected participants to allocate to the Moroccan category a larger percentage of criminal-looking faces than nonmanipulated faces. Moreover, because the trait criminal is both stereotype relevant and negatively valenced, we expected more prejudiced participants to show this effect to a greater extent. Because normative fit should be unaffected by features associated with stereotype-irrelevant traits, we expected participants to allocate to the Moroccan category approximately the same percentage of stupid-looking faces as nonmanipulated faces, despite the fact that the trait stupid is negative and evaluatively congruent with the Moroccan category. Moreover, we expected the allocation of stupid-looking faces to be unaffected by the prejudice of participants.

Method

Participants. Twenty male and 113 female Dutch-speaking students of the Radboud University Nijmegen participated in this study (mean age = 21.67 years, *SD* = 4.73). In return, participants received course credit or €4.

Overview and design. Participants were asked to categorize faces as either Moroccan or non-Moroccan. These faces were Moroccan-looking, criminal Moroccan-looking, and stupid Moroccan-looking. Afterward, implicit prejudice was measured. The experiment employed a mixed-model design with face set (Moroccan-looking vs. criminal Moroccan-looking vs. stupid Moroccan-looking) as a within-subject variable and implicit prejudice as a continuous between-subjects variable. The dependent variable was the percentage of faces categorized as Moroccan.

Materials.

Criminal and stupid features. In order to manipulate criminal-looking facial features, a pilot study was run to construct noisy images of what people thought a typical criminal face looked like. In the pilot, 13 male and 22 female students of the Radboud University Nijmegen (mean age = 21.17 years, *SD* = 2.74) completed a forced-choice version of a reverse-correlation image classification task (Dotsch et al., 2008; Mangini & Biederman, 2004). In this task, participants were repeatedly presented with two stimuli side by side (see Figure 1b). Each stimulus consisted of two random-noise patterns superimposed over a base image (see Figure 1a). The base image in all reverse-correlation tasks throughout this article was the neutral male face of the Averaged Karolinska Directed Emotional Faces Database (Lundqvist & Litton, 1998). The noise patterns were randomly generated at every trial.[2] Within a single trial, one stimulus consisted of the base face with a random-noise pattern and the other with the inverse pattern superimposed. Because the random-noise patterns distorted the base face, the pair of faces appeared to be different in every trial. Participants were instructed to select the most criminal-looking face from each face pair in 770 trials. Averaging all 770 noise patterns selected as most criminal-looking for each participant resulted in classification images representing what our pilot participants thought a typical criminal face looked like (see Figure 2b for the classification image averaged across all participants).

Likewise, in order to manipulate stupid-looking facial features, a pilot study was run to construct noisy images of what people thought a typical stupid face looked like. In this pilot, three male and 28 female students of the Radboud University Nijmegen (mean age = 21.03 years, *SD* = 3.33) completed the same forced-choice reverse-correlation image classification task as in the criminal reverse-correlation pilot but were instructed to select the most stupid-looking face. Trials on which participants responded faster than 300 ms were excluded. Six participants were not included in the resulting averaged classification image, because they responded in less than 300 ms on more than 10% of the trials (this procedure differed slightly from the criminal-looking facial features pilot, because the criminal classification image was generated for an earlier study, reported later in footnote 4). Figure 2c shows the resulting averaged classification image, which represents what participants thought a typical stupid face looked like.

To validate the criminal and stupid classification images, 26 male and 81 female students of the Radboud University Nijmegen (mean age = 22.60 years, *SD* = 4.34) rated all separate 30 criminal and 25 stupid classification images on criminality and stupidity on a 9-point scale ranging from 1 (*not criminal/not stupid*) to 9 (*very criminal/very stupid*). Ratings were averaged across participants for each image. The criminal faces were rated as more criminal ($M = 5.23$, $SD = 0.57$) than the stupid faces ($M = 4.94$, $SD = 0.39$), $t(53) = 2.17$, $p = .03$, $d = 0.59$. The stupid faces were rated as more stupid ($M = 5.02$, $SD = 0.31$) than the criminal faces ($M = 4.55$, $SD = 0.39$), $t(53) = 4.85$, $p < .01$, $d = 1.33$. However, the criminality and stupidity ratings of the stupid faces seemed to be correlated, $r(25) = .36$, $p = .08$. To filter out any characteristics perceived as criminal from the stupid face and vice versa, we orthogonalized the criminality and stupidity ratings and used the resulting values as weights to calculate weighted criminal and stupid classification images (see Figures 2d and 2e). This procedure ensured that changes in perceived criminality by superimposing the weighted criminal classification image did not change perceived stupidity and vice versa. As a result, the orthogonalized criminality and stupidity ratings were uncorrelated, $r(55) = .00$, *ns*, and the parameters underlying the weighted classification images were uncorrelated, $r(4092) = -.04$, *ns*.

[2] The random-noise pattern consisted of superimposed truncated sinusoid images in six orientations (0°, 30°, 60°, 90°, 120°, and 150°) × five spatial frequencies (1, 2, 4, 8, and 16 cycles per image) × two phases (0, $\pi/2$), with random contrasts (amplitudes). In sum, the random noise was a function of 4,092 parameters.

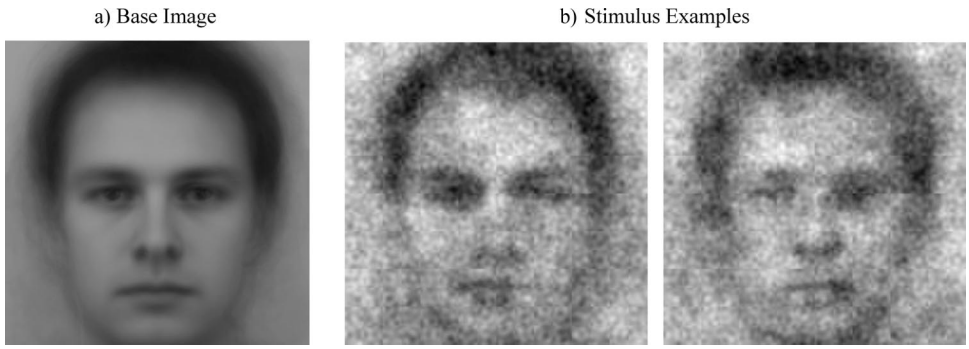

Figure 1. (a) Base image and (b) an example of stimuli used in a trial of the pilot study.

To validate the weighted criminal and stupid classification images, 22 male and 109 female students of the Radboud University Nijmegen (mean age = 21.40 years, $SD = 3.97$) rated the two weighted classification images embedded in a sequence of six other noisy filler images on criminality and stupidity using the same scale as in the previous rating study. The orthogonalization procedure seemed to have improved the extent to which the classification images conveyed the targeted traits: The weighted criminal face was rated as more criminal ($M = 7.76$, $SD = 1.35$) than the weighted stupid face ($M = 5.40$, $SD = 1.74$), $t(130) = 14.54$, $p < .01$, $d = 1.52$. The weighted stupid face was rated as more stupid ($M = 6.05$, $SD = 1.64$) than the weighted criminal face ($M = 5.48$, $SD = 2.06$), $t(130) = 3.25$, $p < .01$, $d = 0.31$.

Stimuli. Three sets of stimuli were used in the experiment: 35 Moroccan-looking, 35 criminal Moroccan-looking, and 35 stupid Moroccan-looking faces. The Moroccan-looking faces were the classification images of the individual participants in the Dotsch et al. (2008, Study 2) experiment (see Figure 2a for the average Moroccan classification image and Figure 3a for an example stimulus). These classification images were the result of the same procedure as in the pilot study described above, but with the instruction for participants to select the most Moroccan-looking face. The criminal and stupid Moroccan-looking faces were generated by, respectively, superimposing the weighted criminal and the weighted stupid classification image on each individual Moroccan-looking face from the Dotsch et al. study. As a result, the final stimuli consisted of 75% Moroccan

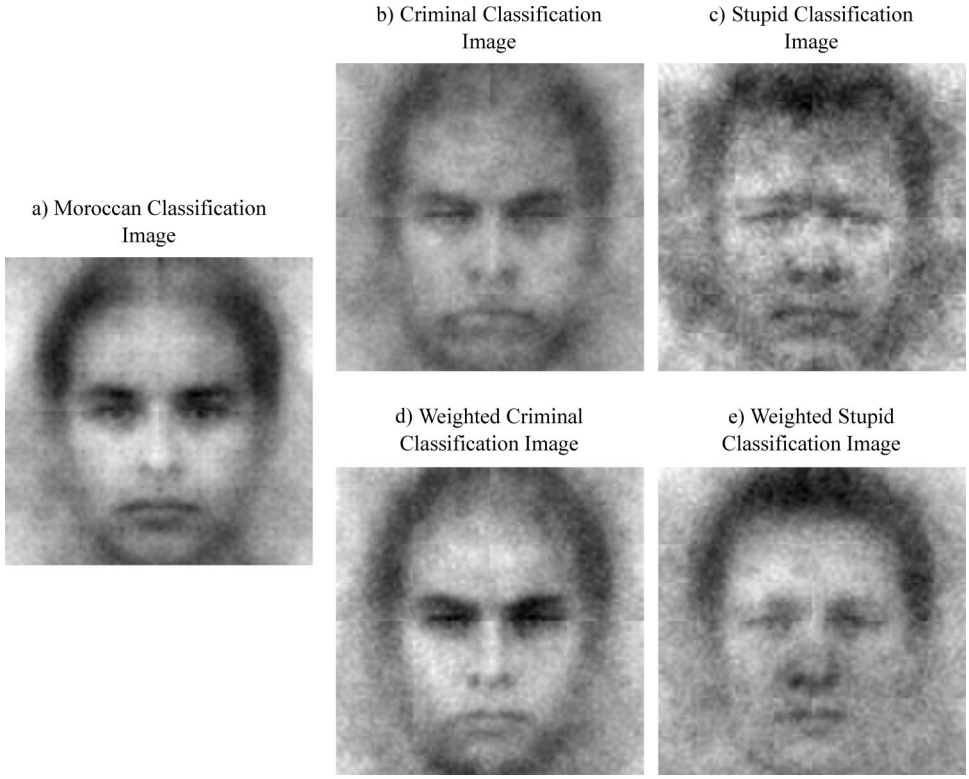

Figure 2. (a) Moroccan classification image, (b) criminal classification image, (c) stupid classification image, (d) weighted criminal classification image, and (e) weighted stupid classification image.

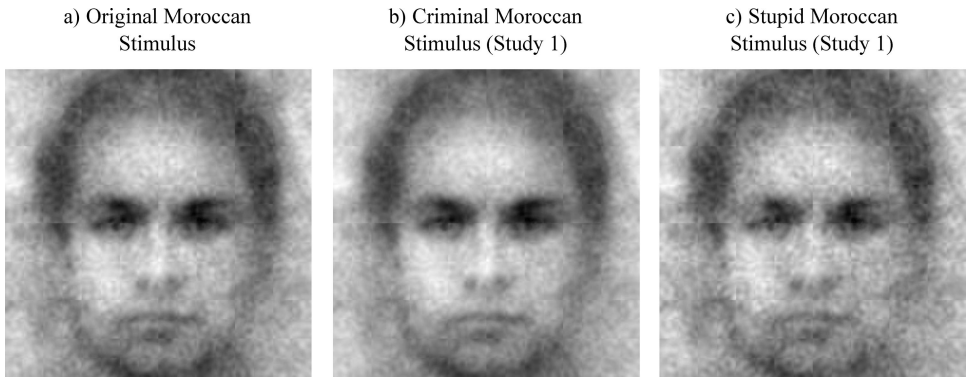

Figure 3. Example stimuli used in Study 1.

classification image + 25% criminal (or stupid) classification image on top of the original base image (see Figures 3b and 3c). As can be seen in Figure 3, the differences between the Moroccan-looking faces and criminal Moroccan-looking faces are subtle and difficult to perceive when presented side by side. A video clip demonstrating the differences more clearly with faces presented in succession is available in the online supplemental materials.

Procedure.

Categorization task. Participants first completed a categorization task in which the 35 Moroccan-looking, the 35 criminal Moroccan-looking, and the 35 stupid Moroccan-looking faces were presented on a computer screen one by one in random order. Each face was presented twice, resulting in a total of 210 trials. Participants were asked to categorize each face as either Moroccan or non-Moroccan as quickly as possible by pressing one of two category keys. They were told that noise was added to the faces to increase task difficulty. Participants were asked to respond as soon as a face appeared on the screen. Faces were presented until a response was given. There was no response window. The intertrial interval was 500 ms.

Implicit prejudice. As a prejudice measure, participants then moved on to a Single-Target Implicit Association Test (ST–IAT; Bluemke & Friese, 2008; Dotsch & Wigboldus, 2008). This task measured indirectly how strongly participants associated Moroccan names (category words, such as *Ibrahim* and *Rachid*) with positive and negative stimuli (valenced words, such as *love*, *peace*, *war*, and *pain*; see the Appendix for the full list of words). We used a shortened version of the ST–IAT, which consisted of practice, compatible, and incompatible blocks.[3] In the practice block, participants classified five positive words with one key and five negative words with another. In the compatible block, participants classified 10 positive words with one key and five negative words and five Moroccan names with another. In the incompatible block, participants classified five positive words and five Moroccan names with one key and 10 negative words with another. Block order was kept constant: The compatible block always preceded the incompatible block. Within blocks, stimuli were presented in random order. When participants made an incorrect classification, error feedback was presented for 1,000 ms. Shorter response latencies on the compatible block than on the incompatible block were assumed to indicate stronger negative than positive associations with Moroccan names (see Greenwald, McGhee, & Schwartz, 1998), which was interpreted as reflecting higher levels of implicit prejudice.

In order to prevent the administration of the ST–IAT to have unwanted effects on the categorization task, we decided on a fixed order of tasks: The categorization task always preceded the ST–IAT.

Results

Implicit prejudice. Incorrect trials and the first two trials of each block were discarded from the ST–IAT analysis. Latencies below 300 ms were set to 300 ms (0.06%). Latencies above 3,000 ms were set to 3,000 ms (0.10%). Analyses were performed on log-transformed latencies, but untransformed mean latencies are reported (in milliseconds). A ST–IAT score was calculated by subtracting the average response latency in the compatible block ($M = 591$, $SD = 103$) from the average response latency in the incompatible block ($M = 627$, $SD = 123$), excluding one participant whose ST–IAT score was lower than 3 *SD* below the mean and three participants who on average responded faster than 300 ms in the categorization task (these four participants were excluded from further analysis). A higher ST–IAT score therefore indicates relatively stronger negative than positive associations with Moroccan names (i.e., being higher in implicit prejudice). On average, participants had stronger negative than positive associations with the category of Moroccans, $t(128) = 5.72$, $p < .01$, $d = 1.01$.

Categorization task. For each participant, the percentage of faces categorized as Moroccan was calculated separately for each face set. Figure 4a plots the results at ± 1 *SD* levels of implicit prejudice. Planned Helmert contrasts in a general linear modeling (GLM) analysis with face set (criminal Moroccan-looking vs. stupid Moroccan-looking vs. Moroccan-looking) as a three-level within-subject factor and implicit prejudice as a continuous factor revealed that participants allocated more criminal Moroccan-looking faces than stupid and original Moroccan-looking faces to the Moroccan category, $F(1, 127) = 95.54$, $p < .01$, $\eta_p^2 = .43$. Additionally, participants allocated fewer stupid Moroccan-looking faces than original Moroccan-looking faces to the Moroccan category, $F(1, 127) = 116.72$, $p < .01$, $\eta_p^2 = .48$. Furthermore, the prejudice slope for criminal Moroccan-looking faces (i.e., the extent to which highly prejudiced participants allocated more criminal Moroccan-looking faces to the Moroccan category than did participants low in prejudice) was greater than the prejudice slope for original Moroccan-looking and stupid Moroccan-looking faces

[3] Throughout the article, the labels *compatible* and *incompatible* reflect the perspective of a prejudiced individual.

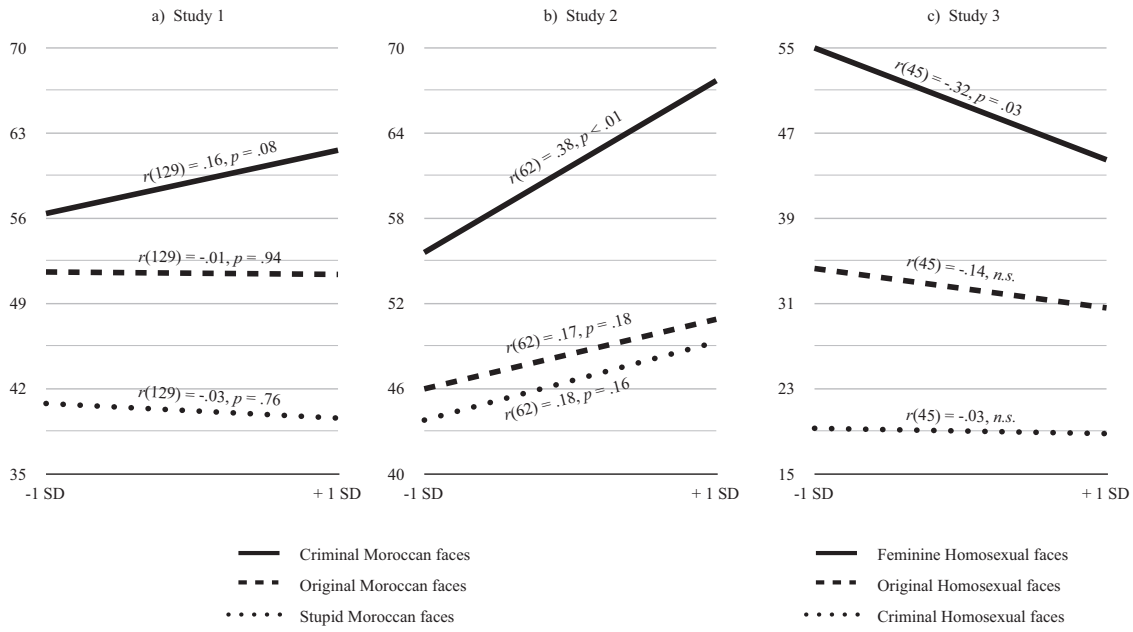

Figure 4. Average percentage of faces allocated to category by face set at ±1 *SD* values of implicit prejudice in (a) Study 1, (b) Study 2, and (c) Study 3.

combined, $F(1, 127) = 4.70$, $p = .03$, $\eta_p^2 = .04$. The prejudice slope for stupid Moroccan-looking faces was not different from the slope for original Moroccan-looking faces, $F(1, 127) < 1$, *ns.* These effects were further explored in two separate GLM analyses with two-level face set factors: one with the levels Moroccan-looking versus criminal Moroccan-looking and one with the levels Moroccan-looking versus stupid Moroccan-looking.

Criminal faces. A GLM analysis with face set (Moroccan-looking vs. criminal Moroccan-looking) as a two-level within-subject factor and implicit prejudice as a continuous factor revealed a main effect of face set, $F(1, 127) = 52.73$, $p < .01$, $\eta_p^2 = .29$, such that participants categorized criminal Moroccan-looking faces ($M = 59.03$, $SD = 16.71$) more often as Moroccan than Moroccan-looking faces ($M = 51.52$, $SD = 16.62$). Additionally, a Face Set × Implicit Prejudice interaction was found, $F(1, 127) = 6.76$, $p = .01$, $\eta_p^2 = .05$, in the expected direction (see Figure 4a): More prejudiced participants overallocated criminal Moroccan-looking faces to the Moroccan category to a greater extent than did less prejudiced participants.

Stupid faces. A GLM analysis with face set (Moroccan-looking vs. stupid Moroccan-looking) as a two-level within-subject factor and implicit prejudice as a continuous factor revealed only a main effect of face set, $F(1, 127) = 116.72$, $p < .01$, $\eta_p^2 = .48$, such that participants categorized stupid Moroccan-looking faces ($M = 40.19$, $SD = 21.02$) less often as Moroccan than Moroccan-looking faces ($M = 51.52$, $SD = 16.62$). No interaction with prejudice was found.

Discussion

As predicted, participants allocated a larger percentage of the criminal-looking faces to the Moroccan category than the nonmanipulated faces. Moreover, more prejudiced participants showed this effect to a greater extent. This bias of highly prejudiced people toward overallocation of criminal-looking faces to the Moroccan category supports our theoretical proposition: Highly prejudiced individuals more strongly endorse the stereotype that Moroccans are criminals and therefore expect Moroccan faces to look more criminal (as has been shown by Dotsch et al., 2008). Because these expectations affect normative fit, people are in general more likely to categorize more criminal-looking faces as Moroccan, but highly prejudiced people do so to a greater extent. This hypothesized effect of prejudice on categorization has been clearly demonstrated in the present study.[4]

Importantly, the effect reported here is different from the effect demonstrated by Dotsch et al. (2008). In Dotsch et al., participants chose the most Moroccan-looking face from two faces with superimposed random noise. On the basis of those choices, a personal classification image could be generated, reflecting the face that would look most Moroccan to each individual participant.

[4] We replicated the effect of overallocation of criminal Moroccan-looking faces and its relation with prejudice in another study that did not include the stupid Moroccan-looking faces and in which the criminal Moroccan-looking faces were created by superimposing the nonorthogonalized criminal classification image. Nineteen male and 73 female Dutch-speaking students of the Radboud University Nijmegen participated in this study (mean age = 21.59 years, $SD = 2.78$). The ST–IAT in this study consisted of double the number of trials as in Study 1, which is more conventional. A GLM analysis with face set (Moroccan-looking vs. criminal Moroccan-looking) as a within-subject factor and implicit prejudice as a continuous factor revealed a main effect of face set, $F(1, 90) = 64.29$, $p < .01$, $\eta_p^2 = .42$, such that participants categorized criminal Moroccan-looking faces ($M = 52.93$, $SD = 19.41$) more often as Moroccan than Moroccan-looking faces ($M = 45.64$, $SD = 18.87$). Additionally a main effect of implicit prejudice, $F(1, 90) = 9.70$, $p < .01$, $\eta_p^2 = .10$, and a Face Set × Implicit Prejudice interaction was found, $F(1, 90) = 5.23$, $p = .03$, $\eta_p^2 = .06$. The means and slopes showed that, in general, participants categorized criminal Moroccan-looking faces more often as Moroccan than Moroccan-looking faces, but more prejudiced participants did so to a greater extent.

Independent participants later rated the classification images of more prejudiced participants to be more criminal than those of less prejudiced participants. In this study, we reversed the procedure, presenting participants with faces constructed on an independent basis to appear more criminal or not and finding that criminal appearance contributed to categorization as Moroccan to a greater extent for more prejudiced individuals. Furthermore, the Dotsch et al. study employed a task in which participants repeatedly compared two faces, a markedly different task from the categorization task employed in this study.

Additionally, participants underallocated stupid faces to the Moroccan category. These results go directly against the evaluative fit hypothesis. Apparently, categorization is affected by more than valence alone. Both criminality and stupidity are negative traits and therefore congruent with the negative valence of the Moroccan category. However, only superimposing features associated with the stereotype-relevant trait criminal on Moroccan faces elicited overallocation. On the other hand, superimposing features associated with the stereotype-irrelevant trait stupid evoked underallocation. Thus, valence contributes to fit only to the extent that it derives from stereotype-relevant traits. Moreover, more prejudiced participants did not allocate more stupid Moroccan-looking faces to the Moroccan category. In fact, we did not find any relation between prejudice and the allocation of stupid Moroccan-looking faces to the Moroccan category. This suggests that the influence of prejudice on category allocation is also restricted to faces with features associated with stereotype-relevant traits.

We did not predict the underallocation of stupid faces. Because the trait stupid was assumed to be stereotype irrelevant, we expected that superimposing features associated with stupidity would not matter for Moroccan face categorization and that faces with those features would therefore have approximately the same probability to be categorized as Moroccan as the original faces. However, it should be noted that to test for the effect of two independent trait dimensions (criminal and stupid) on the allocation of Moroccan faces, we orthogonalized the dimensions. In doing so, we might have compromised the meaning of the facial features superimposed on the Moroccan classification images. It is likely that the verbal trait labels (*stupid*, *criminal*) were no longer represented optimally by the resulting orthogonalized classification images, which precludes any hard conclusions about the underlying responsible factors. We therefore set out to run a new study in which we used nonorthogonalized classification images.

The observed correlation between ST–IAT scores and allocation of criminal Moroccan-looking faces to the Moroccan category was predicted but might have resulted from the particular set of stimuli we selected for the ST–IAT. One might argue that our set of valenced ST–IAT words might to some extent also reflect stereotypical Moroccan traits. Specifically, words such as *love*, *peace*, *safe*, *emergency*, or *pain* might be indirectly related to the trait criminal (although, please note that words such as *aggressive* or *criminal* were not part of the stimulus set). This offers an alternative account of why ST–IAT scores predicted the allocation of criminal—but not stupid—Moroccan-looking faces to the Moroccan category.

Although evidence for the influence of stimulus features on the magnitude of the IAT effect is mixed (Bluemke & Friese, 2006; De Houwer, 2001; Mitchell, Nosek, & Banaji, 2003), some studies demonstrated that different sets of valenced words do not necessarily produce equal IAT effects (e.g., Steffens & Plewe, 2001). To rule out the possibility that the correlation between ST–IAT scores and the observed categorization pattern was the result of a possible confound of valence with stereotypicality in the ST–IAT valenced words, we used a different set of valenced words, unrelated to the trait criminal, in our next study.

Study 2

Study 2 was aimed at replicating the results of Study 1, while ruling out two alternative explanations. First, as argued above, the underallocation of stupid Moroccan-looking faces might be a result of our orthogonalization procedure. We therefore generated faces using the nonorthogonalized criminal and stupid classification images. Second, in Study 1, ST–IAT scores predicted the allocation of criminal—but not stupid—Moroccan-looking faces to the Moroccan category. As argued above, the observed correlations might have been the result of a possible confound of valence with stereotypicality in the ST–IAT valenced words. In Study 2, we therefore manipulated the ST–IAT valenced words to either be the same as in Study 1 (original stimulus set) or to be completely unconfounded with stereotypicality (new stimulus set). Moreover, in Study 1 we used an unconventional shortened version of the ST–IAT. In Study 2 we instead used an ST–IAT of conventional length. Apart from these changes, the study was identical to Study 1.

Despite that the criminal and stupid classification images were not orthogonalized, we expected that participants would again overallocate the criminal Moroccan-looking faces. Moreover, we expected participants to no longer underallocate the stupid Moroccan-looking faces, as we suspected that the orthogonalization procedure was responsible for the underallocation.

Because features of IAT valenced words typically do not influence IAT effect magnitude (e.g., De Houwer, 2001), we expected that regardless of the specific set of valenced words used in the ST–IAT (original vs. new set), ST–IAT scores would predict the overallocation of criminal Moroccan-looking faces. Because the trait stupid is stereotype irrelevant, features associated with that trait should not affect normative fit, despite evaluative congruency. We therefore expected the allocation of stupid Moroccan-looking faces to be unrelated to ST–IAT scores, regardless of the specific set of valenced words used.

Method

Participants. Four male and 60 female Dutch-speaking students of the Radboud University Nijmegen participated in this study (mean age = 19.58 years, $SD = 2.39$). In return, participants received course credit.

Overview and design. Participants were asked to categorize faces as either Moroccan or non-Moroccan. These faces were Moroccan-looking, criminal Moroccan-looking, and stupid Moroccan-looking. Afterward, implicit prejudice was measured using an ST–IAT. The experiment used a 3 (face set: Moroccan-looking vs. criminal Moroccan-looking vs. stupid Moroccan-looking) × 2 (ST–IAT stimulus set: original vs. new) × Implicit Prejudice mixed-model design, with face set as a within-subject variable, ST–IAT stimulus set as a between-subjects variable, and

implicit prejudice as a continuous between-subjects variable. The dependent variable was the percentage of faces categorized as Moroccan.

Materials.

Categorization task stimuli. Three sets of stimuli were used in the experiment: 35 Moroccan-looking, 35 criminal Moroccan-looking, and 35 stupid Moroccan-looking faces. The Moroccan-looking faces were the same as in Study 1 (see Figure 5a for an example). The criminal and stupid Moroccan-looking faces were generated by, respectively, superimposing the nonorthogonalized criminal and stupid classification image resulting from the pilot study in Study 1 on each individual Moroccan-looking face (see Figures 5b and 5c). Again, the differences between the three classes of stimuli were very subtle.

Procedure. The procedure of Study 2 was the same as the procedure of Study 1, except for where indicated below. Participants started with the categorization task using the set of stimuli generated on the basis of the nonorthogonalized criminal and stupid classification images. As a prejudice measure, participants then moved on to the same ST–IAT used in Study 1. However, the ST–IAT this time consisted of twice the number of trials per block. Participants were randomly assigned to one of two ST–IAT stimulus set conditions. In the original stimulus set condition, the valenced words were identical to the valenced words in Study 1. In the new stimulus set condition, all valenced words that might be indirectly related to the (for Moroccans) stereotypical trait criminal (i.e., *death*, *emergency*, *loss*, *pain*, *love*, *peace*, *safe*, and *friend*) were replaced by valenced words unrelated to that trait (i.e., *accident*, *rotten*, *poison*, *stink*, *vomit*, *laugh*, *fun*, *gift*, and *vacation*). See the Appendix for the full list of stimuli used.

Results

Implicit prejudice. Incorrect trials and the first two trials of each block were discarded from the ST–IAT analysis. There were no latencies below 300 ms or above 3,000 ms in the data. Analyses were performed on log-transformed latencies, but untransformed mean latencies are reported (in milliseconds). An ST–IAT score was calculated by subtracting the average response latency in the compatible block ($M = 587$, $SD = 75$) from the average response latency in the incompatible block ($M = 631$, $SD = 94$), excluding one participant whose ST–IAT score was higher than 3 *SD* above the mean and one participant who on average responded faster than 300 ms in the categorization task (these two participants were excluded from further analysis). Again, a higher ST–IAT score indicates relatively stronger negative than positive associations with Moroccan names (i.e., being higher in implicit prejudice). On average, participants had stronger negative than positive associations with the category of Moroccans, $t(61) = 4.53$, $p < .01$, $d = 1.16$.

Categorization task. For each participant, the percentage of faces categorized as Moroccan was calculated separately for each face set. We expected that participants would allocate more of the criminal-looking faces than the original Moroccan-looking faces to the Moroccan category, while allocating approximately the same percentage of stupid-looking faces to the Moroccan category as the original Moroccan-looking faces. Moreover, we expected that participants' prejudice would predict the percentage of criminal—but not stupid or original—Moroccan-looking faces they allocated to the Moroccan category. We also expected that these effects would be observed in the original ST–IAT stimulus set condition ($n = 32$) as well as in the new ST–IAT stimulus set condition ($n = 30$).

Figure 4b plots the results at ± 1 *SD* levels of implicit prejudice. Planned Helmert contrasts in a 3 (face set: criminal Moroccan-looking vs. stupid Moroccan-looking vs. Moroccan-looking) $\times$ 2 (ST–IAT stimulus set: original vs. new) $\times$ Implicit Prejudice GLM analysis, with face set as a within-subject factor, ST–IAT stimulus set as a between-subjects factor, and implicit prejudice as a continuous between-subjects factor, revealed that participants allocated more criminal Moroccan-looking faces than stupid and original Moroccan-looking faces to the Moroccan category, $F(1, 58) = 77.63$, $p < .01$, $\eta_p^2 = .57$. There was no difference between the stupid and original Moroccan-looking faces in allocation to the Moroccan category, $F(1, 58) = 2.89$, *ns.* Furthermore, the prejudice slope for criminal Moroccan-looking faces (i.e., the extent to which highly prejudiced participants allocated more criminal Moroccan-looking faces to the Moroccan category than did participants low in prejudice) was greater than the prejudice slope for original Moroccan-looking and stupid Moroccan-looking faces combined, $F(1, 58) = 4.41$, $p < .04$, $\eta_p^2 = .05$. The prejudice slope for stupid Moroccan-looking faces was not different from the slope for original Moroccan-looking faces, $F(1, 58) < 1$, *ns.* Importantly, none of these effects were moderated by ST–IAT stimulus set (all Fs < 1, *ns*). We therefore dropped the ST–IAT stimulus set variable in subsequent analyses. The observed effects were further

Figure 5. Example stimuli used in Study 2.

explored in two separate GLM analyses with two-level face set factors: one with the levels original Moroccan-looking versus criminal Moroccan-looking and one with the levels original Moroccan-looking versus stupid Moroccan-looking.

Criminal faces. A GLM analysis with face set (Moroccan-looking vs. criminal Moroccan-looking) as a two-level within-subject factor and implicit prejudice as a continuous between-subjects factor replicated the main effect of face set found in Study 1, $F(1, 60) = 62.66$, $p < .01$, $\eta_p^2 = .51$, such that participants categorized criminal Moroccan-looking faces ($M = 60.81$, $SD = 17.89$) more often as Moroccan than Moroccan-looking faces ($M = 47.81$, $SD = 17.61$). Additionally, we found a main effect of implicit prejudice, $F(1, 60) = 5.75$, $p = .02$, $\eta_p^2 = .09$, such that more prejudiced participants categorized more faces as Moroccan. Most important, we replicated the Face Set × Implicit Prejudice interaction from Study 1, $F(1, 60) = 4.97$, $p = .03$, $\eta_p^2 = .08$, in the expected direction (see Figure 4b): More prejudiced participants overallocated criminal Moroccan-looking faces to the Moroccan category to a greater extent than did less prejudiced participants.

Stupid faces. A GLM analysis with face set (Moroccan-looking vs. stupid Moroccan-looking) as a two-level within-subject factor and implicit prejudice as a continuous between-subjects factor on the percentage of faces categorized as Moroccan revealed a marginally significant face set effect, $F(1, 60) = 3.27$, $p = .08$, $\eta_p^2 = .05$, such that participants allocated fewer stupid Moroccan-looking faces ($M = 45.88$, $SD = 18.76$) than original Moroccan-looking faces ($M = 47.81$, $SD = 17.61$) to the Moroccan category. No interaction with prejudice was found.

Discussion

With Study 2 we replicated the findings of Study 1, showing that participants overallocated criminal Moroccan-looking faces to the Moroccan category and that this overallocation of criminal Moroccan-looking faces was greater for more prejudiced participants. As in Study 1, prejudice moderated the overallocation of criminal Moroccan-looking faces to the Moroccan category, whereas prejudice did not affect the allocation of stupid Moroccan-looking faces to the Moroccan category.

The valenced words used in the ST–IAT to measure prejudice in Study 1 were valenced but potentially confounded with the Moroccan stereotype. The words could indirectly be linked to the stereotypical trait criminal. In Study 2 we manipulated the valenced words to be either the same as in Study 1 or to be stereotype unconfounded. This manipulation did not moderate any of the effects reported in Study 2, suggesting that in our case the selection of specific valenced words did not influence the relationship between prejudice and category allocation.

Although the effect was marginally significant, we again observed underallocation of stupid Moroccan-looking faces to the Moroccan category in Study 2. This suggests that, although orthogonalization may have been responsible for a certain amount of underallocation, there might be other causes.

One potential cause is that stupidity might not be stereotype irrelevant but counterstereotypical. However, it is very unlikely that the trait stupid is counterstereotypical for Moroccans. When asked to name stereotypical Moroccan traits in the Gordijn et al. (2001) study, participants should have named traits opposite to stupid (e.g., smart) for stupid to qualify as counterstereotypical. They did not. Moreover, when we asked 17 male and 93 female students at Radboud University Nijmegen (mean age = 21.46 years, $SD = 4.66$) to judge to what extent the average Dutch person believes Moroccans to be stupid on a 7-point scale ranging from 1 (*not at all*) to 7 (*very much*), participants on average responded with 4.73 ($SD = 1.13$). Although this response was significantly lower than the extent to which our participants thought that the average Dutch person believes Moroccans to be criminal ($M = 5.39$, $SD = 0.90$), $t(109) = 6.34$, $p < .01$, it was still above the midpoint of the scale. This seems to indicate that it is unlikely that stupidity is counterstereotypical for Moroccans, although we cannot fully rule out this possibility on the basis of the current data.

We believe that two potential mechanisms may underlie the underallocation of faces with features associated with the trait stupid. First, traits that are stereotype irrelevant for one category (e.g., Moroccans) can be stereotype relevant for other accessible categories, thereby increasing the likelihood of a non-Moroccan response, despite the trait not having any relevance to the Moroccan stereotype. Second, superimposing features associated with stereotype-irrelevant traits might simply, by mere dilution, reduce the proportion of typically Moroccan facial features in the stimulus face. Under this interpretation, features associated with stereotype-irrelevant traits add nonmeaningful noise to the signal, thereby decreasing the likelihood of a Moroccan response. Both processes could work in tandem to make underallocation of faces with facial features associated with stereotype-irrelevant traits likely. Furthermore, these processes explain why we did not find a relation between prejudice toward Moroccans and underallocation of faces with facial features associated with the stereotype-irrelevant trait stupid: The Moroccan stereotype—and therefore prejudice toward Moroccans—is not involved in the processes causing underallocation.

Taken together, Studies 1 and 2 reliably demonstrated that social categorization is biased. Subtle differences in facial features associated with stereotype-relevant traits caused overallocation of faces possessing those features. The extent to which faces with those features were overallocated was predicted by participants' prejudice: More negatively prejudiced participants overallocated faces with features associated with a negative stereotype-relevant trait to a greater extent. This did not happen for just any negative trait. The allocation of faces with features associated with a negative stereotype-irrelevant trait was unrelated to participants' level of prejudice. These results provide strong evidence against a pure evaluative fit hypothesis and fully support our normative fit hypothesis.

Study 3

In both Studies 1 and 2, the trait criminal enhanced allocation of faces to the Moroccan category. Moreover, these studies demonstrated that not just any negative trait enhances overallocation to a stigmatized category (as would be predicted by the evaluative fit hypothesis), but only negative stereotype-relevant traits. In Study 3 we aimed to test whether the same negative trait that elicits overallocation to one stigmatized category would not elicit overallocation for another stigmatized category. Moreover, Study 3 included a positive stereotype-relevant trait, which, if it still causes

overallocation of faces conveying that trait for a stigmatized category, would be even stronger evidence against the evaluative fit hypothesis and in favor of a normative fit explanation.

Study 3 employed the category of homosexual men, which can be perceived from faces (Rule & Ambady, 2008; Rule, Ambady, Adams, & Macrae, 2008). Two traits were used in this experiment: the trait criminal—which is stereotype irrelevant for homosexual men in the Netherlands—and the stereotype-relevant trait feminine. Femininity is the most stereotypic personality trait ascribed to homosexual men (Madon, 1997). We expected participants to overallocate feminine homosexual-looking faces to the homosexual category, but not criminal homosexual-looking faces. Because the previous studies showed that the influence of prejudice is restricted to stereotype-relevant traits, we also expected that if any effect of prejudice on category allocation would emerge, it would be obtained only for faces with features associated with the stereotype-relevant trait, that is, feminine.

Whether femininity is a positive trait has been debated in the literature. In most—but not all—empirical work, femininity or stereotypically female traits have been found to be more positive than masculine traits (e.g., Der-Karabetian & Smith, 1977; Silvern & Ryan, 1983; but see Eagly & Mladinic, 1994). Within the context of faces, a meta-analysis on attractiveness (Rhodes, 2006) showed that feminine female faces are clearly more attractive. More important for the current study, there is evidence that people have a preference for feminized male faces (Penton-Voak, Jacobson, & Trivers, 2004; Perrett et al., 1998; Rhodes, Hickford, & Jeffery, 2000), particularly when female facial averages are used to transform male facial images (Rennels, Bronstad, & Langlois, 2008; Rhodes, 2006). This is similar to the procedure in this study: superimposing a feminine classification image onto a male base face. In a pilot study (reported in the Method section below), we show that adding feminine features, in the context of our stimuli, indeed caused male faces to be evaluated more positively.

Because femininity, in the context of the current stimuli, is a positive trait, we expected participants who were negatively prejudiced toward the homosexual category to allocate fewer of the feminine faces to the homosexual category. We expected prejudice toward homosexuals to be unrelated to the allocation of faces with features associated with the stereotype-irrelevant trait criminal.

Method

Participants. Three male and 42 female heterosexual Dutch-speaking students of the Radboud University Nijmegen participated in this study (mean age = 21.76 years, $SD = 2.52$). In return, participants received course credit or €4.

Overview and design. Participants were asked to categorize faces as either homosexual or nonhomosexual. These faces were homosexual-looking, criminal homosexual-looking, and feminine homosexual-looking. Afterward, implicit prejudice was measured. The experiment used a mixed-model design with face set (homosexual-looking vs. criminal homosexual-looking vs. feminine homosexual-looking) as a within-subject variable and implicit prejudice toward homosexuality as a continuous between-subjects variable. The dependent variable was the percentage of faces categorized as homosexual.

Materials.

Homosexual and feminine features. In order to create homosexual-looking faces and to manipulate feminine-looking facial features, a pilot study was run to construct noisy images of what people thought typical homosexual and feminine faces looked like. In the pilot, 13 male and 54 female students of the Radboud University Nijmegen (mean age = 22.19 years, $SD = 5.25$) completed the same forced-choice reverse-correlation image classification task as in the previous pilots, with the same male base image. Thirty-five participants were instructed to select the most homosexual-looking face, whereas 32 other participants were instructed to select the most feminine-looking face. Trials on which participants responded faster than 300 ms were removed. Figure 6 shows the resulting averaged classification images, which represent what participants thought typical homosexual and feminine faces looked like.

Stimuli. Three sets of stimuli were used in the experiment: 35 homosexual-looking, 35 criminal homosexual-looking, and 35 feminine homosexual-looking faces. The nonorthogonalized criminal classification image generated in the pilot of Study 1 was used to generate the criminal homosexual-looking faces. The criminal and feminine homosexual-looking faces were generated from the individual homosexual-looking classification images in the same way as the stimuli in Studies 1 and 2 (see Figure 6 for example stimuli), that is, by superimposing the averaged criminal or averaged feminine classification images, respectively, on each individual homosexual-looking classification image. Again, the resulting differences between the three classes of stimuli were very subtle. All stimuli used the same male base face as in the previous two studies.

To ensure that the feminine homosexual-looking faces were indeed more positive than the original and criminal homosexual-looking faces, 11 male and 76 female students from Radboud University Nijmegen (mean age = 20.71 years, $SD = 3.03$) rated all 105 stimuli on valence using a 9-point scale ranging from –4 (*very negative*) to 4 (*very positive*). Indeed, the feminine homosexual-looking faces were rated more positive ($M = 0.47$, $SD = 0.66$) than the original homosexual-looking faces ($M = 0.19$, $SD = 0.62$), $t(86) = 7.30$, $p < .01$, $d = 1.57$, which were rated more positive than the criminal homosexual-looking faces ($M = -0.66$, $SD = 0.74$), $t(86) = 17.24$, $p < .01$, $d = 3.72$.

Procedure. The procedure of Study 3 was the same as the procedure of Study 2, except for where indicated below.

Categorization task. Participants categorized one by one and in random order 35 homosexual-looking, 35 criminal homosexual-looking, and 35 feminine homosexual-looking faces into homosexual and nonhomosexual categories. There were two blocks, and each face was presented once in each block, resulting in a total of 210 trials.

Implicit prejudice. As a prejudice measure, participants then moved on to a homosexual–heterosexual IAT (Banse, Seise, & Zerbes, 2001; Steffens & Buchner, 2003). This task measured indirectly how strongly participants associated male homosexual couples (e.g., John + Richard) and heterosexual couples (e.g., Matthew + Julia) with positive and negative words (see the Appendix for all stimuli). The IAT consisted of one compatible, one incompatible, and three practice blocks, which were always

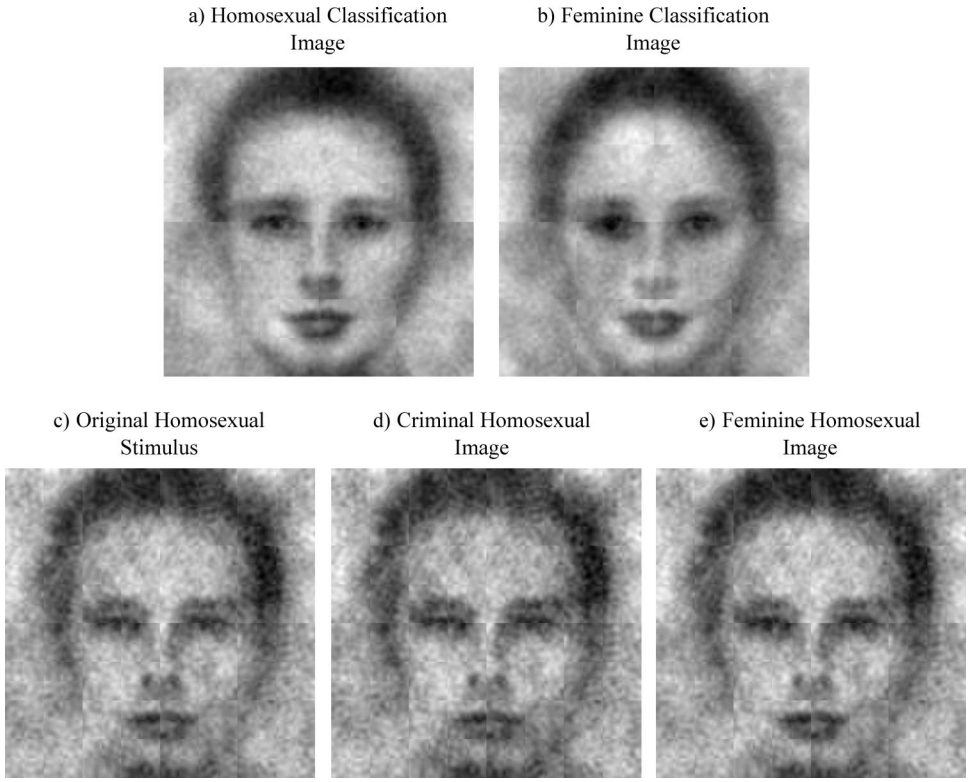

Figure 6. (a) Homosexual classification image, (b) feminine classification image, and (c, d, e) example stimuli used in Study 3.

carried out in the following order: In the first practice block, participants classified 10 positive words with the left key and 10 negative words with the right. In the compatible practice block, participants classified 10 heterosexual couples with the left key and 10 homosexual couples with the right. In the following critical compatible block, participants classified 10 positive words and 10 heterosexual couples with the left key and 10 negative words and 10 homosexual couples with the right key. The incompatible practice and critical blocks followed, with the functions of the keys with respect to the couples switched. Shorter response latencies on the compatible block than on the incompatible block were assumed to indicate stronger negative than positive associations with homosexual couples compared with heterosexual couples, which was interpreted as reflecting higher levels of implicit prejudice.

Results

Implicit prejudice. Incorrect trials and the first two trials of each block were discarded from the IAT analysis. There were no latencies below 300 ms. Latencies above 3,000 ms were set to 3,000 ms (0.32%). Analyses were performed on log-transformed latencies, but untransformed mean latencies are reported (in milliseconds). An IAT score was calculated by subtracting the average response latency in the compatible block ($M = 710$, $SD = 122$) from the average response latency in the incompatible block ($M = 830$, $SD = 214$). A higher IAT score therefore indicates relatively stronger negative than positive associations with homosexual couples compared with heterosexual couples (i.e., being higher in implicit prejudice). Participants on average had stronger negative than positive associations with the homosexual category relative to the heterosexual category, $t(45) = 5.65$ $p < .01$, $d = 0.69$; that is, the category of homosexuals was a stigmatized category on the implicit level for this sample.

Categorization task. For each participant the percentage of faces categorized as homosexual was calculated separately for each face set. Figure 4c plots the results at ± 1 *SD* levels of implicit prejudice. Planned Helmert contrasts in a GLM analysis with face set (feminine homosexual-looking vs. criminal homosexual-looking vs. homosexual-looking) as a three-level within-subject factor and implicit prejudice as a continuous factor revealed that participants allocated more feminine homosexual-looking faces than criminal and original homosexual-looking faces to the homosexual category, $F(1, 43) = 171.67$, $p < .01$, $\eta_p^2 = .80$. Participants also allocated fewer criminal than original homosexual-looking faces to the homosexual category, $F(1, 43) = 122.32$, $p < .01$, $\eta_p^2 = .74$. Furthermore, the prejudice slope for feminine homosexual-looking faces (i.e., the extent to which highly prejudiced participants allocated fewer feminine homosexual-looking faces to the homosexual category than did participants low in prejudice) was greater than the prejudice slope for original homosexual-looking and criminal homosexual-looking faces combined, $F(1, 43) = 5.17$, $p = .03$, $\eta_p^2 = .11$. The prejudice slope for criminal homosexual-looking faces did not differ from the slope for original homosexual-looking faces, $F(1, 43) = 1.69$, *ns.* These effects were further explored in two separate GLM analyses with

two-level face set factors: one with the levels homosexual-looking versus feminine homosexual-looking faces and one with the levels homosexual-looking versus criminal homosexual-looking faces.

Feminine faces. A GLM analysis with face set (homosexual-looking vs. feminine homosexual-looking) as a within-subject factor and implicit prejudice as a continuous factor revealed the expected main effect of face set, $F(1, 43) = 101.26$, $p < .01$, $\eta_p^2 = .70$, such that participants categorized feminine homosexual-looking faces ($M = 49.75$, $SD = 16.29$) more often as homosexual than homosexual-looking faces ($M = 32.48$, $SD = 13.77$). Additionally, a predicted Face Set × Implicit Prejudice interaction was marginally significant, $F(1, 43) = 3.86$, $p = .06$, $\eta_p^2 = .08$: More prejudiced participants overallocated feminine homosexual-looking faces to a lesser extent to the homosexual category (see Figure 4c).

Criminal faces. A GLM analysis with face set (homosexual-looking vs. criminal homosexual-looking) as a within-subject factor and implicit prejudice as a continuous factor revealed only a main effect of face set, $F(1, 130) = 122.32$, $p < .01$, $\eta_p^2 = .74$, such that participants categorized criminal homosexual-looking faces ($M = 19.05$, $SD = 10.93$) less often as homosexual than homosexual-looking faces ($M = 32.48$, $SD = 13.77$). As predicted, we found no interaction with prejudice.

Discussion

As expected, participants overallocated feminine faces, but not criminal faces, to the homosexual category. These results demonstrate that the categorization of faces from different stigmatized categories relies on different trait dimensions conveyed by the faces: Whereas the task of categorizing faces as Moroccan in Studies 1 and 2 elicited overallocation of criminal faces, the task of categorizing faces as homosexual did not elicit overallocation of criminal faces. Thus, faces with features associated with negative traits are not necessarily more likely to be allocated to a negatively evaluated category. This is a strong argument in favor of the stereotype specificity of normative fit and against the pure evaluative fit position.

The observed underallocation of criminal homosexual-looking faces can be explained by the same mechanisms that were invoked to explain the underallocation of stupid Moroccan-looking faces to the Moroccan category (see the Discussion of Study 2). That is, although criminality may be irrelevant to the stereotype of homosexuals, it may be a very salient relevant trait to the stereotypes of other categories. Facial features associated with the trait criminal therefore increase the likelihood of a nonhomosexual response and, hence, result in underallocation to the homosexual category. Also, superimposing a stereotype-irrelevant trait like criminality may dilute the proportion of homosexual features in the stimulus face, which may also have contributed to the observed underallocation. Because both these presumed processes (i.e., higher fit with other categories and dilution of typical traits) are unrelated to prejudice toward homosexuals, a relationship with prejudice was neither expected nor observed.

Additionally, implicit prejudice toward homosexuals moderated the allocation of the feminine faces, such that more prejudiced participants showed less overallocation of feminine homosexual-looking faces. This finding sheds more light on the way normative fit affects categorization. Faces that convey stereotype-relevant evaluative congruence have better fit for highly prejudiced individuals. In the case of the homosexual category, the negativity associated with criminality derives from a trait that is not stereotype relevant and therefore does not provide better fit for prejudiced individuals. Femininity, on the other hand, is stereotype relevant, but positively valenced. Because prejudiced individuals evaluate the category as more negative, faces conveying less of the stereotype-relevant positive trait feminine have better fit. Alternatively, using a less traditional definition of prejudice as an evaluative continuum ranging from positive to negative, one could also interpret this result as follows: Because more positively prejudiced individuals evaluate the category of homosexuals as more positive, faces conveying more of the stereotype-relevant positive trait feminine have better fit. This interpretation might seem less counterintuitive. Nonetheless, under both interpretations, prejudice affects only the allocation of faces with features associated with stereotype-relevant traits.

General Discussion

In three studies we have shown that social categorization, a process from which so many biases originate, is biased itself, at the level of category allocation. Across studies, we have shown a double dissociation. In Studies 1 and 2, participants overallocated faces conveying the stereotype-relevant negative trait criminal to the Moroccan category, especially if they were prejudiced. The stereotype-irrelevant negative trait stupid, on the other hand, did not elicit overallocation. In Study 3, using the stigmatized category homosexual, the previously used negative trait criminal did not elicit overallocation, but the stereotype-relevant positive trait femininity did. These results demonstrate that normative fit is enhanced for faces conveying stereotype-relevant traits, regardless of valence. Moreover, individual differences in implicit prejudice predicted the extent to which stereotype-relevant traits caused overallocation: Whereas more negatively prejudiced people showed greater overallocation of faces conveying negative stereotype-relevant traits, they showed less overallocation of faces conveying positive stereotype-relevant traits. These results strongly support our normative fit hypothesis: There is better fit for faces with features that are evaluatively congruent on the condition that those features are associated with stereotype-relevant traits. More prejudiced individuals have a stronger negative evaluation of the category and therefore show enhanced normative fit for negative stereotype-relevant traits and less normative fit for positive stereotype-relevant traits.

Our results are especially striking considering the very subtle manipulation of criminal, stupid, and feminine features. Looking at the example stimuli in Figures 3 or 5, it is difficult to see any difference between the original Moroccan and the criminal-looking Moroccan images. Nonetheless, our subtle manipulation, based on pilot participants' criminality judgments, was strong enough to elicit a difference in category allocation between highly and less prejudiced participants. Our findings converge with the results of recent studies by, for example, Blair, Judd, Sadler, and Jenkins (2002) and Livingston and Brewer (2002) in underscoring the importance of considering facial features in the study of prejudice and discrimination.

Another strength of the current work is the use of reverse-correlation methods to construct Moroccan-, homosexual-,

criminal-, stupid-, and feminine-looking faces. At no point in the current research were any subjective notions of the researchers about what, for example, a criminal face looks like able to influence the results. Bottom-up data-driven research methodologies are gaining traction in research on higher level cognition (e.g., Dotsch et al., 2008; Langner, Becker, & Rinck, 2009; Oosterhof & Todorov, 2008) and have as a primary advantage that we are testing what is in participants' minds, not what is in researchers' minds (see Mangini & Biederman, 2004; Todorov, Dotsch, Wigboldus, & Said, under review). It is not what the researchers believe a criminal face looks like that biases social categorization but what participants believe a criminal face looks like.

Our work utilized a paradigm in which participants were forced to decide whether to allocate faces to one specific accessible category. However, social psychologists have theorized that categories are activated automatically on perception of category members (Allport, 1954; Devine, 1989; Dovidio, Evans, & Tyler, 1986; for an overview, see Macrae & Bodenhausen, 2000). In daily life, people are mostly not on the lookout for members of one specific category, as in our task. The cognitive system has ample choice of social categories to activate (e.g., Moroccan, homosexual, Dutch, Black, Mexican, female, social psychologist). Future research should address the extent to which the effects obtained generalize to spontaneous categorization given multiple categories to allocate to.

Our results are in line with findings within the emotion perception literature that racially ambiguous faces with an angry expression are more likely to be categorized as Black by more prejudiced people (Hugenberg & Bodenhausen, 2004; Hutchings & Haddock, 2008; also see Bijlstra, Holland, & Wigboldus, 2010). On the basis of the former, Hugenberg and Sacco (2008, p. 1057) suggested that "when attributes of a social target are consistent with the attributes of a particular social category (e.g., angry expressions with the Black stereotype), this can strongly influence category selection." The current work extends previous work by stating more precisely the effect of prejudice and stereotype content on social categorization and by demonstrating that category selection might sometimes be biased by prejudice in the opposite direction: For the homosexual category, more negatively prejudiced individuals were less prone to overallocate faces conveying the positive stereotype-relevant trait feminine.

The current work further advances knowledge about the relationship between prejudice and stereotypes. As noted by Wittenbrink et al. (1997), the definition of prejudice as an individual's negative attitude toward an outgroup is widely accepted (Ashmore, 1970; Dovidio & Gaertner, 1986; Stroebe & Insko, 1989). A stereotype, on the other hand, reflects culturally shared beliefs (Devine, 1989) or personal beliefs about an outgroup (Krueger, 1996). Amodio and Devine (2006) provided evidence for the independence of stereotyping and prejudice. On the other hand, Wittenbrink et al. (1997, p. 271) concluded that on the implicit level, stereotyping and prejudice are intertwined. We go beyond both conclusions by demonstrating in what manner stereotyping and prejudice might be intertwined: The biased allocation by prejudiced people depends on stereotype-relevant valence. In this view, the stereotype constrains the effects of prejudice.

The demonstrated biases in category allocation might function as both a stereotype-maintenance and a prejudice-maintenance device (as has been alluded to by Dotsch et al., 2008, p. 980). The reported experiments showed a general tendency to bias categorization in the direction of a culturally shared stereotype, thereby supporting the propagation of a cultural meme (Dawkins, 1976): As feminine-looking homosexuals are more likely to be categorized as homosexual, the likelihood that a counterstereotypical (e.g., less feminine-looking) exemplar will be categorized as homosexual becomes smaller. Without counterstereotypical category members, there is no need to change the stereotype. Moreover, the reported experiments showed a specific tendency of prejudice to bias categorization in the direction of evaluative congruency: As prejudiced individuals are more likely to categorize more negative and fewer positive exemplars into stigmatized categories, the proportion of evaluatively incongruent exemplars allocated to stigmatized categories becomes smaller. If prejudiced people encounter relatively more negative members of stigmatized groups (because they are more prone to categorize them as such), there is less reason to change their evaluation of those groups. They remain prejudiced.

Social categorization (Allport, 1954) has proven to be an important concept in explaining prejudice (see Fiske, 2005). Ample research has shown the influence of prejudice as a consequence of categorization. The current research clearly demonstrates the influence of prejudice on the categorization process itself.

References

Allport, G. W. (1954). *The nature of prejudice.* Reading, MA: Addison-Wesley.

Allport, G. W., & Kramer, B. M. (1946). Some roots of prejudice. *Journal of Psychology, 22,* 9–39. doi:10.1080/00223980.1946.9917293

Amodio, D. M., & Devine, P. G. (2006). Stereotyping and evaluation in implicit race bias: Evidence for independent constructs and unique effects on behavior. *Journal of Personality and Social Psychology, 91,* 652–661. doi:10.1037/0022-3514.91.4.652

Ashmore, R. D. (1970). Prejudice: Causes and cures. In B. E. Collins (Ed.), *Social psychology: Social influence, attitude change, group processes and prejudice* (pp. 245–339). Reading, MA: Addison-Wesley.

Banse, R., Seise, J., & Zerbes, N. (2001). Implicit attitudes towards homosexuality: Reliability, validity, and controllability of the IAT. *Zeitschrift für Experimentelle Psychologie, 48,* 145–160.

Bijlstra, G., Holland, R. W., & Wigboldus, D. H. J. (2010). The social face of emotion recognition: Evaluations versus stereotypes. *Journal of Experimental Social Psychology, 46,* 657–663. doi:10.1016/j.jesp.2010.03.006

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. doi:10.1037/0022-3514.83.1.5

Bluemke, M., & Friese, M. (2006). Do features of stimuli influence IAT effects? *Journal of Experimental Social Psychology, 42,* 163–176. doi:10.1016/j.jesp.2005.03.004

Bluemke, M., & Friese, M. (2008). Reliability and validity of the Single-Target IAT (ST–IAT): Assessing automatic affect towards multiple attitude objects. *European Journal of Social Psychology, 38,* 977–997. doi:10.1002/ejsp.487

Bruner, J. S. (1957). On perceptual readiness. *Psychological Review, 64,* 123–152. doi:10.1037/h0043805

Castano, E., Yzerbyt, V., Bourguignon, D., & Seron, E. (2002). Who may enter? The impact of in-group identification on in-group/out-group categorization. *Journal of Experimental Social Psychology, 38,* 315–322. doi:10.1006/jesp.2001.1512

Coenders, M., Lubbers, M., Scheepers, P., & Verkuyten, M. (2008). More

than two decades of changing ethnic attitudes in the Netherlands. *Journal of Social Issues, 64,* 269–285. doi:10.1111/j.1540-4560.2008.00561.x

Dawkins, R. (1976). *The selfish gene.* Oxford, England: Oxford University Press.

De Houwer, J. (2001). A structural and process analysis of the Implicit Association Test. *Journal of Experimental Social Psychology, 37,* 443–451. doi:10.1006/jesp.2000.1464

Der-Karabetian, A., & Smith, A. J. (1977). Sex-role stereotyping in the United States: Is it changing? *Sex Roles, 3,* 193–198. doi:10.1007/BF00288670

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56,* 5–18. doi:10.1037/0022-3514.56.1.5

Dotsch, R., & Wigboldus, D. H. J. (2008). Virtual prejudice. *Journal of Experimental Social Psychology, 44,* 1194–1198. doi:10.1016/j.jesp.2008.03.003

Dotsch, R., Wigboldus, D. H. J., Langner, O., & van Knippenberg, A. (2008). Ethnic out-group faces are biased in the prejudiced mind. *Psychological Science, 19,* 978–980. doi:10.1111/j.1467-9280.2008.02186.x

Dovidio, J. F., Evans, N., & Tyler, R. B. (1986). Racial stereotypes: The contents of their cognitive representations. *Journal of Experimental Social Psychology, 22,* 22–37. doi:10.1016/0022-1031(86)90039-9

Dovidio, J. F., & Gaertner, S. L. (1986). Prejudice, discrimination, and racism: Historical trends and contemporary approaches. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 1–34). Orlando, FL: Academic Press.

Dubow, S. (1992). Afrikaner nationalism, apartheid, and the conceptualization of "race." *The Journal of African History, 33,* 209–237. doi:10.1017/S0021853700032217

Eagly, A. H., & Mladinic, A. (1994). Are people prejudiced against women? Some answers from research on attitudes, gender stereotypes, and judgements of competence. *European Review of Social Psychology, 5,* 1–35. doi:10.1080/14792779543000002

Elliott, D. N., & Wittenberg, B. H. (1955). Accuracy of identification of Jewish and non-Jewish photographs. *Journal of Abnormal and Social Psychology, 51,* 339–341. doi:10.1037/h0043566

Fiske, S. T. (1998). Stereotyping, prejudice, and discrimination. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (4th ed., Vol. 2, pp. 357–411). New York, NY: McGraw-Hill.

Fiske, S. T. (2005). Social cognition and the normality of prejudgment. In J. F. Dovidio, P. Glick, & L. A. Rudman (Eds.), *On the nature of prejudice: Fifty years after Allport* (pp. 36–53). Malden, MA: Blackwell.

Gordijn, E. H., Koomen, W., & Stapel, D. A. (2001). Level of prejudice in relation to knowledge of cultural stereotypes. *Journal of Experimental Social Psychology, 37,* 150–157. doi:10.1006/jesp.2000.1443

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. doi:10.1037/0022-3514.74.6.1464

Himmelfarb, S. (1966). Studies in the perception of ethnic group members: I. Accuracy, response bias, and anti-Semitism. *Journal of Personality and Social Psychology, 4,* 347–355. doi:10.1037/h0023808

Hugenberg, K., & Bodenhausen, G. V. (2004). Ambiguity in social categorization: The role of prejudice and facial affect in race categorization. *Psychological Science, 15,* 342–345. doi:10.1111/j.0956-7976.2004.00680.x

Hugenberg, K., & Sacco, D. F. (2008). Social categorization and stereotyping: How social categorization biased person perception and face memory. *Social and Personality Psychology Compass, 2,* 1052–1072. doi:10.1111/j.1751-9004.2008.00090.x

Hutchings, P. B., & Haddock, G. (2008). Look Black in anger: The role of implicit prejudice in the categorization and perceived emotional intensity of racially ambiguous faces. *Journal of Experimental Social Psychology, 44,* 1418–1420. doi:10.1016/j.jesp.2008.05.002

Kosic, A., & Phalet, K. (2006). Ethnic categorization of immigrants: The role of prejudice, perceived acculturation strategies and group size. *International Journal of Intercultural Relations, 30,* 769–782. doi:10.1016/j.ijintrel.2006.06.003

Krueger, J. (1996). Personal beliefs and cultural stereotypes about racial characteristics. *Journal of Personality and Social Psychology, 71,* 536–548. doi:10.1037/0022-3514.71.3.536

Langner, O., Becker, E. S., & Rinck, M. (2009). Social anxiety and anger identification: Bubbles reveal differential use of facial information with low spatial frequencies. *Psychological Science, 20,* 666–670. doi:10.1111/j.1467-9280.2009.02357.x

Le Pelley, M. E., Reimers, S. J., Calvini, G., Spears, R., & Beesley, T. (2010). Stereotype formation: Biased by association. *Journal of Experimental Psychology: General, 139,* 138–161. doi:10.1037/a0018210

Livingston, R. W., & Brewer, M. B. (2002). What are we really priming? Cue-based versus category-based processing of facial stimuli. *Journal of Personality and Social Psychology, 82,* 5–18. doi:10.1037/0022-3514.82.1.5

Lundqvist, D., & Litton, J. E. (1998). *The Averaged Karolinska Directed Emotional Faces: AKDEF* [CD-ROM]. Stockholm, Sweden: Karolinska Institutet.

Macrae, C. N., & Bodenhausen, G. V. (2000). Social cognition: Thinking categorically about others. *Annual Review of Psychology, 51,* 93–120. doi:10.1146/annurev.psych.51.1.93

Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Calvini, G. (1999). Seeing more than we can know: Visual attention and category activation. *Journal of Experimental Social Psychology, 35,* 590–602. doi:10.1006/jesp.1999.1396

Macrae, C. N., Bodenhausen, G. V., Milne, A. B., 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. doi:10.1006/jesp.1997.1328

Madon, S. (1997). What do people believe about gay males? A study of stereotype content and strength. *Sex Roles, 37,* 663–685. doi:10.1007/BF02936334

Mangini, M. C., & Biederman, I. (2004). Making the ineffable explicit: Estimating the information employed for face classifications. *Cognitive Science, 28,* 209–226. doi:10.1207/s15516709cog2802_4

Mitchell, J. P., Nosek, B. A., & Banaji, M. R. (2003). Contextual variations in implicit evaluation. *Journal of Experimental Psychology: General, 132,* 455–469. doi:10.1037/0096-3445.132.3.455

Oakes, P. J. (1987). The salience of social categories. In J. C. Turner, M. A. Hogg, P. J. Oakes, S. D. Reicher, & M. S. Wetherell (Eds.), *Rediscovering the social group: A self-categorization theory* (pp. 117–141). Oxford, England: Blackwell.

Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences, USA, 105,* 11087–11092. doi:10.1073/pnas.0805664105

Penton-Voak, I. S., Jacobson, A., & Trivers, R. (2004). Population differences in attractiveness judgments of male and female faces: Comparing British and Jamaican samples. *Evolution and Human Behavior, 25,* 355–370. doi:10.1016/j.evolhumbehav.2004.06.002

Perrett, D. I., Lee, K. J., Penton-Voak, I., Rowland, D., Yoshikawa, S., Burt, D. M., . . . Akamatsu, S. (1998, August 27). Effects of sexual dimorphism on facial attractiveness. *Nature, 394,* 884–887. doi:10.1038/29772

Pettigrew, T. F., Allport, G. W., & Barnett, E. O. (1958). Binocular resolution and perception of race in South Africa. *British Journal of Psychology, 49,* 265–278.

Quanty, N. B., Keats, J. A., & Harkins, S. G. (1975). Prejudice and criteria

for identification of ethnic photographs. *Journal of Personality and Social Psychology, 32,* 449–454. doi:10.1037/h0077093

Rennels, J. L., Bronstad, P. M., & Langlois, J. H. (2008). Are attractive men's faces masculine or feminine? The importance of type of facial stimuli. *Journal of Experimental Psychology: Human Perception and Performance, 34,* 884–893. doi:10.1037/0096-1523.34.4.884

Rhodes, G. (2006). The evolutionary psychology of facial beauty. *Annual Reviews of Psychology, 57,* 199–226. doi:10.1146/annurev.psych.57.102904.190208

Rhodes, G., Hickford, C., & Jeffery, L. (2000). Sex-typicality and attractiveness: Are supermale and superfemale faces super-attractive? *British Journal of Psychology, 91,* 125–140. doi:10.1348/000712600161718

Rule, N. O., & Ambady, N. (2008). Brief exposures: Male sexual orientation is accurately perceived at 50 ms. *Journal of Experimental Social Psychology, 44,* 1100–1105. doi:10.1016/j.jesp.2007.12.001

Rule, N. O., Ambady, N., Adams, R. B., Jr., & Macrae, C. N. (2008). Accuracy and awareness in the perception and categorization of male sexual orientation. *Journal of Personality and Social Psychology, 95,* 1019–1028. doi:10.1037/a0013194

Ruys, K. I., Dijksterhuis, A., & Corneille, O. (2008). On the (mis)categorization of unattractive brides and attractive prostitutes. *Experimental Psychology, 55,* 182–188. doi:10.1027/1618-3169.55.3.182

Silvern, L. E., & Ryan, V. L. (1983). A reexamination of masculine and feminine sex-role ideals and conflicts among ideals for the man, woman, and person. *Sex Roles, 9,* 1223–1248. doi:10.1007/BF00303104

Stangor, C., & Schaller, M. (1996). Stereotypes as individual and collective representations. In C. N. Macrae, C. Stangor, & M. Hewstone (Eds.), *Stereotypes and stereotyping* (pp. 3–40). New York, NY: Guilford Press.

Steffens, M. C., & Buchner, A. (2003). Implicit Association Test: Separating transsituationally stable and variable components of attitudes toward gay men. *Experimental Psychology, 50,* 33–48. doi:10.1027//1618-3169.50.1.33

Steffens, M. C., & Plewe, I. (2001). Items' cross-category associations as a confounding factor in the Implicit Association Test. *Zeitschrift für Experimentelle Psychologie, 48,* 123–134.

Stroebe, W., & Insko, C. A. (1989). Stereotypes, prejudice, and discrimination: Changing conceptions in theory and research. In D. Bar-Tal, C. F. Grauman, A. W. Kruglanski, & W. Stroebe (Eds.), *Stereotypes and prejudice: Changing conceptions* (pp. 3–34). New York, NY: Springer.

Tajfel, H. (1969). Cognitive aspects of prejudice. *Journal of Social Issues, 25,* 79–97. doi:10.1111/j.1540-4560.1969.tb00620.x

Todorov, A., Dotsch, R., Wigboldus, D. H. J., & Said, C. P. (under review). Data-driven methods for modeling social perception. *Social and Personality Psychology Compass.*

Turner, J. C. (1987). A self-categorization theory. In J. C. Turner, M. A. Hogg, P. J. Oakes, S. D. Reicher, & M. S. Wetherell (Eds.), *Rediscovering the social group: A self-categorization theory* (pp. 42–67). Oxford, England: Blackwell.

van Knippenberg, A., & Dijksterhuis, A. (2000). Social categorization and stereotyping: A functional perspective. In W. Stroebe & M. Hewstone (Eds.), *European Review of Social Psychology* (Vol. 11, pp. 105–144). West Sussex, England: Wiley.

van Knippenberg, A., van Twuyver, M., & Pepels, J. (1994). Factors affecting social categorization processes in memory. *British Journal of Social Psychology, 33,* 419–431.

Verkuyten, M., & Zaremba, K. (2005). Interethnic relations in a changing political context. *Social Psychology Quarterly, 68,* 375–386. doi:10.1177/019027250506800405

Wittenbrink, B., Judd, C. M., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology, 72,* 262–274. doi:10.1037/0022-3514.72.2.262

Zebrowitz, L. A. (1996). Physical appearance as a basis of stereotyping. In C. N. Macrae, C. Stangor, & M. Hewstone (Eds.), *Stereotypes and stereotyping* (pp. 79–120). New York, NY: Guilford Press.

(*Appendix follows*)

Appendix

Stimuli From Studies 1, 2, and 3

Stimuli Used in the ST–IAT in Study 1

Moroccan Names

Ibrahim, *Mustafa*, *Abdul*, *Aziz*, *Tarik*, *Mohammed*, *Achmed*, *Youssef*, *Rachid*, *Adil*

Negative Words

Accident (*ongeluk*), *death* (*dood*), *sadness* (*verdriet*), *emergency* (*nood*), *loss* (*verlies*), *pain* (*pijn*), *disgust* (*walging*), *sickness* (*ziekte*), *bitterness* (*bitter*), *cancer* (*kanker*)

Positive Words

Love (*liefde*), *peace* (*vrede*), *safe* (*veilig*), *healthy* (*gezond*), *nice* (*leuk*), *luck* (*geluk*), *happy* (*blij*), *friend* (*vriend*), *pretty* (*mooi*), *party* (*feest*)

Stimuli Used in the ST–IAT in Study 2

Moroccan Names

Ibrahim, *Mustafa*, *Abdul*, *Aziz*, *Tarik*, *Mohammed, Achmed*, *Youssef*, *Rachid*, *Adil*

Negative Words

Accident (*ongeluk*), *rotten* (*verrot*), *sadness* (*verdriet*), *poison* (*vergif*), *stink* (*stank*), *vomit* (*braaksel*), *disgust* (*walging*), *sickness* (*ziekte*), *bitterness* (*bitter*), *cancer* (*kanker*)

Positive Words

Laugh (*lachen*), *fun* (*plezier*), *gift* (*kado*), *healthy* (*gezond*), *nice* (*leuk*), *luck* (*geluk*), *happy* (*blij*), *vacation* (*vakantie*), *pretty* (*mooi*), *party* (*feest*)

Stimuli Used in the IAT in Study 3

Heterosexual Couples

Bart + Fleur, *Gijs + Femke*, *Daniel + Anne*, *Bram + Emma*, *Luuk + Anouk*

Fleur + Bart, *Femke + Gijs*, *Anne + Daniel*, *Emma + Bram*, *Anouk + Luuk*

Homosexual Couples

Jan + Kees, *Pieter + Thijs*, *Sander + Jeroen*, *Maarten + Frank*, *Joost + Michiel*

Kees + Jan, *Thijs + Pieter*, *Jeroen + Sander*, *Frank + Maarten*, *Michiel + Joost*

Negative Words

Accident (*ongeluk*), *death* (*dood*), *loss* (*verlies*), *pain* (*pijn*), *disgust* (*walging*), *mad* (*kwaad*), *sickness* (*ziekte*), *unpleasant* (*vervelend*), *mean* (*gemeen*), *danger* (*gevaar*)

Positive Words

Peace (*vrede*), *safe* (*veilig*), *healthy* (*gezond*), *funny* (*leuk*), *nice* (*aardig*), *luck* (*geluk*), *win* (*winnen*), *enjoyment* (*vermaak*), *pretty* (*mooi*), *pleasurable* (*prettig*)

Received April 14, 2010
Revision received December 26, 2010
Accepted December 29, 2010 ■