European Journal of Social Psychology, *Eur. J. Soc. Psychol.* (2012)
Published online in Wiley Online Library (wileyonlinelibrary.com) **DOI**: 10.1002/ejsp.1928

Research article

Behavioral information biases the expected facial appearance of members of novel groups

RON DOTSCH*, DANIEL H. J. WIGBOLDUS AND AD VAN KNIPPENBERG
Behavioural Science Institute, Radboud University Nijmegen, The Netherlands

Abstract

The present study tests the hypothesis that behavioral information diagnostic of an out-group's traits biases the expected facial appearance of out-group members toward having facial features corresponding with the inferred traits. Participants formed a stereotype about a novel group based on random exemplar faces, presented alongside descriptions of their behavior. The behavioral information was manipulated to reflect either trustworthy or criminal traits, whereas the stimulus faces did not reflect any traits. Afterwards, participants' expected facial appearance of group members was assessed using a reverse correlation task. Independent judges rated the resulting visualized expectations as more criminal in the criminal behavioral information condition than in the trustworthy behavioral information condition. The current work establishes a causal link between behavioral information and expected out-group faces where previously only correlations had been observed. Copyright © 2012 John Wiley & Sons, Ltd.

People form a great number of stereotypes in the course of their lives, through processes of (non-)conscious covariance detection (e.g., Ford & Stangor, 1992; Jussim, 1991; Lewicki, 1986), illusory correlation (Hamilton & Gifford, 1976; Hamilton & Sherman, 1989; Mullen & Johnson, 1990), or social learning and communication (e.g., Hirschfeld, 1995; Wigboldus, Semin, & Spears, 2000). These formation processes can be simulated in lab experiments (e.g., Crawford, Sherman, & Hamilton, 2002; Sherman, 1996; Smith & Zarate, 1990). In these experiments, participants typically receive information about traits or behaviors of exemplars of a novel category. If those exemplars, on average, possess a specific trait and variability across exemplars is low, associations between groups and traits will be formed (Dijksterhuis & van Knippenberg, 1999), especially when group members are perceived as a unified entity (i.e., highly entitative; Campbell, 1958; Crawford et al., 2002; Hamilton & Sherman, 1996).

Most stereotype formation research assessed the impact of newly formed stereotypes at the level of verbal trait and behavior descriptions. On the basis of these measurements, one might assume stereotype content to be abstract and schematic (e.g., Higgins & Bargh, 1987). However, stereotypes were originally defined as "pictures in our head" (Lippmann, 1922). McArthur and Baron (1983) emphasized the contribution of physical appearance to stereotyping, as did Brewer (1988), Carlston (1994), and McGarty (2002). Perceptual information may even be the primary basis of conceptual knowledge (Barsalou, 1999). In this view, the nature of stereotype content is for an important part perceptual (Zebrowitz, 1996). Perceptual stereotype content might contain visual information (e.g., group members' typical physical appearance) as well as information from any other modality. For instance, the auditive component of the stereotype might consist of what people expect typical group members' speech to sound like, in terms of, for example, typical pitch or accented speech (e.g., Gluszek & Dovidio, 2010).

Being multi-modal structures themselves, stereotypes influence processes on non-semantic levels, such as face perception. Specifically, faces are automatically categorized by matching bottom-up perceptual information with top-down stereotypical knowledge about categories (e.g., Freeman & Ambady, 2011). For instance, stereotypical knowledge about the category of homosexuals provides a template for the kind of low-level visual input that is perceived as a gay face (Dotsch, Wigboldus, & van Knippenberg, 2011). Indeed, ample studies have demonstrated how existing categories and stereotypes bias social and emotional face perception (e.g., Bijlstra, Holland, & Wigboldus, 2010; Dotsch et al., 2011; Eberhardt, Dasgupta, & Banaszynski, 2003; Hugenberg, 2005; Hugenberg & Bodenhausen, 2004; Hutchings & Haddock, 2008), but none have investigated the influence of *newly* formed stereotypes on face perception. Here, we focus on the influence of newly formed stereotypes on the expected facial appearance of out-group members.

One could argue that the expected facial appearance of members of a group is simply the aggregate of all previously encountered exemplar faces. In line with Dijksterhuis and van Knippenberg (1999), the stereotypical facial appearance associated with a specific group would consist of the average face configuration that exemplars of the group display with

*Correspondence to: Ron Dotsch, Behavioural Science Institute, Radboud University Nijmegen, PO Box 9104, 6500 HE Nijmegen, The Netherlands.
E-mail: r.dotsch@bsi.ru.nl

low variability. However, faces are not just pure feature configurations but take on social meaning as people infer traits from faces (e.g., Todorov, Said, Engell, & Oosterhof, 2008; Zebrowitz & Montepare, 2008). This opens up the possibility that encountered exemplar faces might contribute to the inference of stereotypical traits. Perhaps even more interestingly, the reverse might also take place: other sources than the face, such as exemplar behavior, may inform stereotypical facial appearance. Indeed, people use knowledge about someone's personality as a source of information about an individual's facial appearance (Hassin & Trope, 2000). We suggest that this finding should generalize from the individual to the group level, that is, that people use knowledge about exemplar group members' personality as a source of information about the expected facial appearance of group members.

Assessing Expected Facial Appearance

One promising technique to assess the expected facial appearance of group members is the so-called reverse correlation method (Dotsch & Todorov, 2011; Gosselin & Schyns, 2003; Todorov, Dotsch, Wigboldus, & Said, 2011). Reverse correlation is a data-driven method that allows researchers to probe internal representations of categories without making any a priori assumptions about what those internal representations might look like. The output of the task is an image that visualizes people's criteria for classifying a face as belonging to a specific target category. Reverse correlation tasks have been used to visualize the expected facial appearance of various out-groups (e.g., Dotsch, Wigboldus, Langner, & van Knippenberg, 2008; Dotsch et al., 2011; Imhoff, Dotsch, Bianchi, Banse, & Wigboldus, 2011). For instance, Dotsch et al. (2008) visualized what Dutch participants expected typical Moroccan faces to look like. Moroccans are a stigmatized immigrant minority group in the Netherlands (see Coenders, Lubbers, Scheepers, & Verkuyten, 2008) and are strongly associated with the trait criminal (Gordijn, Koomen, & Stapel, 2001). In the task, participants were instructed to select from two faces, shown side-by-side, the face that according to them was most Moroccan-looking. In fact, all faces in the task consisted of the same (non-Moroccan) base face but with different random noise patterns superimposed. Because the random noise distorted the base face, every face had a different appearance. The average composite of the noise patterns that participants, across a large number of trials, classified as Moroccan is called the classification image. When superimposed on the original base face, the classification image visualized what people think a typical Moroccan face looks like. These classification images were rated by independent judges, and these ratings revealed that, on average, participants expected Moroccan faces to look criminal. Moreover, classification images of participants who had a more negative attitude toward Moroccans, that is, who were more prejudiced, had an even more criminal appearance. The reverse correlation technique provides an excellent means for assessing participants' subjective expectations of the typical facial appearance of group members without forcing participants into being biased: The presented faces were pure random variations (without any a priori manipulations that may communicate the researchers' intention), and the word criminal was never mentioned in the task. Therefore, any criminal features in the Dotsch et al. classification images reflected spontaneous use of criminal facial features in performing a Moroccan face reverse correlation task.

The Current Work

In the Dotsch et al. (2008) reverse correlation task, the criminal features in the expected facial appearance of Moroccans might have been a result of negative traits inferred about the group, or they might be a consequence of the actual or perceived facial configurations of Moroccans in daily life. The use of real-world groups, as in the Dotsch et al. paper, precluded drawing any causal conclusions, because participants came to the lab with stereotypes (i.e., expected facial appearance and typical traits) ready formed. It is therefore unclear whether, as argued earlier, traits inferred from the behavior of previously encountered exemplars informed people's expectations about the expected facial appearance of group members. Here, we propose that behavioral information diagnostic of an out-group's traits will bias perceivers' expectations of typical out-group facial appearance toward having facial features corresponding with the inferred traits, even if the original out-group faces did not have those features.

In order to establish for the first time this causal link, we aimed to experimentally demonstrate that the facial representation of a novel group is affected by available behavioral information about its exemplars. Trustworthy behavioral information should cause people to expect faces of a novel group to look more trustworthy, whereas criminal behavioral information should cause people to expect faces of a novel group to look more criminal. We aimed to provide a conservative test of this hypothesis by providing participants with more pertinent information than the behavioral information: the actual faces of the group members. The facial representation of the novel group could therefore resemble the mere aggregate of the presented faces of the group members, ignoring any available behavioral information.

We tested this idea using a stereotype formation task in which we showed descriptions of exemplar behavior alongside exemplar faces of the novel group X. We manipulated exemplar behavioral information to reflect either a trustworthy group X or a criminal group X. Importantly, between behavioral information conditions, group X faces did not differ. Subsequently, we assessed participants' expected group X facial appearance using a reverse correlation task similar to Dotsch et al. (2008), in which participants repeatedly selected the face that according to them was most likely to be a group X member. We predicted that the classification images (depicting participants' expected group X facial appearance) would vary in line with the behavioral information manipulation. Note that this does not necessarily have to be the case: the reverse correlation task leaves participants opportunity to not use the behavioral information, as participants could in principle rely solely on the appearance of the original group X exemplar faces when performing the reverse correlation task.

Moreover, in line with the findings of Dotsch et al. (2008) that prejudice (i.e., a negative attitude) predicted bias in the classification images, we measured participants' explicit and implicit attitude toward the novel group. Although implicit attitudes are often thought to stem from long-term socialization experiences (e.g., Rudman, 2004; Rudman, Phelan, & Heppen, 2007; but see Castelli, Carraro, Gawronski, & Gava, 2010), both explicit and implicit attitudes toward groups can be induced in the lab in just a single stereotype formation session (see Ratliff & Nosek, 2010). When no norm exists that prevents participants from explicitly stating a negative evaluation, there usually is a correlation between explicit and implicit attitudes (Greenwald, Poehlman, Uhlmann, & Banaji, 2009; also see Gawronski & Bodenhausen, 2006). We therefore expected both explicit and implicit attitudes to vary in line with the behavioral information manipulation and to correlate with classification image bias.

METHOD

Participants

Seventy-seven students (13 men, $M_{age} = 22.36$, $SD = 4.91$) from Radboud University Nijmegen participated in the stereotype formation study. They received €5 or course credits for participating.

Overview and Design

The basic experiment consisted of four tasks, the stereotype formation task, the reverse correlation task, an Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998) and explicit measures. In the stereotype formation task, participants formed an impression of two novel groups, X and Y, based on exemplar faces and behavioral descriptions. For half of the participants, the group X behaviors were predominantly indicative of the trait *trustworthy* and group Y behaviors predominantly indicative of the trait *criminal*. For the other half, this was reversed. Note that we were interested in the induced group X stereotype only, as group X was the only group being visualized in the reverse correlation task, but we included group Y in the formation task to offset the behaviors of group X. Importantly, in all cases, the exemplar faces did not differ between experimental behavioral information conditions. After a filler task, we visualized participants' expectations of typical group X faces using a reverse correlation task (Dotsch et al., 2008; Mangini & Biederman, 2004), yielding so-called classification images. To assess the extent to which participants' group X classification images varied in line with the trustworthy or criminal behavior manipulation, a new group of independent participants rated each participant's resulting classification image on both traits.

The experimental design consisted of one between-participants factor (group X behavioral information: trustworthy versus criminal). Additionally, the design included several counter-balancing factors to control for, for example, effects of order and exemplar faces. These will be explained in more detail later. After the reverse correlation task, we measured implicit and explicit attitude toward group X.[1] We decided to use this order for every participant, because we thought it more likely that the IAT would influence participants' responses in the reverse correlation task than that the reverse correlation task would influence responses in the IAT. After all, people may become aware of their implicit bias while responding in the IAT (De Houwer, Teige-Mocigemba, Spruyt, & Moors, 2009; Monteith, Voils, & Ashburn-Nardo, 2001), whereas this, presumably, is not the case in the reverse correlation task.

Materials

Stereotype Formation Task

In the stereotype formation task (adapted from Crawford et al., 2002), participants were instructed to form an impression of two novel groups, X and Y, based on exemplar faces and behavioral descriptions presented on a computer screen. They received this information in two blocks, one block per group. Block order was counter-balanced. In total, participants viewed 20 exemplar faces and behaviors per group. In each trial, a label of the current group ("Group X" or "Group Y") was presented at the top of the screen, with an exemplar face in the center of the screen, and a behavioral description beneath it. The task was self-paced, that is, participants could take as much time as they wished to look at each face-behavior pair.

The exemplar faces consisted of one of two base faces (one for each group, counter-balanced across participants, see Figure 1(a)) with superimposed noise that distorted the faces to create subtle variations on the base face. The two base faces were selected from the Radboud Face Database (Langner et al., 2010). The selected faces (models 7 and 23) were matched in valence and showed a neutral expression. The images were cropped to 512×512 pixels, converted to grayscale, and blurred with a low-pass Gaussian filter (with a kernel spanning 20 pixels in both image axes) to better match the spatial frequency band of the noise. The noise was constructed in the same manner as in Dotsch et al. (2008), that is, by superimposing multiple layers of sinusoid patches with random amplitudes in six orientations, five spatial frequencies, and two phases.[2] We created two sets of 20 noise patterns, one set for each base face (counter-balanced across participants), resulting in 20 exemplar faces for one group and 20 exemplar faces for the other group (see Figure 1 (b) for examples). Thus, each group was associated with a set of faces that appeared to have more variability between groups than within groups. Please note that as a result of this procedure, the exemplar faces differed from each other only on the basis of random noise patterns and the two different base faces.

Exemplar faces of one group were paired with 10 neutral behaviors (e.g., "This member of group X/Y crosses the street") and 10 trustworthy behaviors (e.g., "This member of group X/Y returns the wallet he found"). Exemplar faces of the other group were paired with 10 different neutral behaviors and 10 criminal behaviors (e.g., "This member of group X/Y

[1]After this procedure, we also administered a Moroccan IAT, a Dutch Need for Closure scale (Cratylus, 1995) and a Need for Cognition scale (Cacioppo & Petty, 1982) to explore potentially relevant individual differences, but these measures proved unrelated to the processes investigated and are therefore not reported in this article.

[2]The random noise pattern consisted of superimposed truncated sinusoid images in 6 orientations (0º, 30º, 60º, 90º, 120º, and 150º) × 5 spatial frequencies (2, 4, 8, 16, and 32 cycles per image) × 2 phases (0, π/2), with random contrasts (amplitudes). In sum, the random noise was a function of 4092 parameters, each defining the contrast value of one truncated sinusoid spanning two cycles.

Figure 1. Base faces (a) and example stimuli in stereotype formation task for each base image (b)

robs another person in an alley"). The specific face-behavior pairings were randomized within groups. Which set of 10 neutral behaviors was associated with which group was counter-balanced across participants. Exemplar order was randomized.

Reverse Correlation Task

To visualize the face participants expected typical group X members to have, we used a two-image forced choice reverse correlation task adapted from Dotsch et al. (2008; also see Dotsch & Todorov, 2011; Mangini & Biederman, 2004). In 480 trials, participants repeatedly chose from two faces presented side-by-side the face that most likely was a member of group X. All faces in this task were generated using as base face the average of the two blurred base faces from the stereotype formation task (Figure 2(a)). Averaging the two base faces ensured that stimuli in the reverse correlation task were equally likely to resemble group X and group Y faces. We generated random noise patterns in exactly the same manner as for the stereotype formation task (Footnote 1) and Dotsch et al. (2008). In a single trial, one stimulus consisted of the base face with a random noise pattern and the other with the negative pattern superimposed (Figure 2(b)). Because the random noise patterns distorted the base face, the faces appeared to be different in every trial. All participants received identical sets of noisy stimuli but in random order. Averaging all noise patterns chosen as group X member per participant resulted in individual classification images, representing what a participant expected typical group X members to look like.[3]

Implicit Association Test

In the IAT, participants classified positive (e.g., love) and negative words (e.g., death), and the group labels "Group X" and "Group Y" into categories. Our IAT consisted of five blocks in the following fixed order: (1) a practice block of 20 trials in which participants classified positive words with the left key and negative words with the right key; (2) a practice block of 20 trials in which participants classified positive words and the label "Group X" with the left key and negative words and the label "Group Y" with the right key; (3) a target block of 40 trials, with the same key mapping as the previous block; (4) a practice block of 20 trials, in which participants practiced the reversed mapping of the group labels: classifying positive words and the label "Group Y" with the left key and negative words and the label "Group X" with the right key; (5) a target block of 40 trials, with the same key mapping as the previous block. We chose to present the blocks in fixed order, because (i) we were interested in variance due to our manipulation only; keeping the order constant allowed us to reduce error, and (ii) to keep the number of counter-balancing factors in our study limited.

In each block, every type of stimulus was shown an equal number of times. Within blocks, the order of stimuli was randomized. Faster responses in Blocks 2 and 3 than in Blocks 4 and 5 were assumed to indicate stronger positive than negative associations with group X and stronger negative than positive associations with group Y. We computed IAT scores using the D-measure algorithm (Greenwald, Nosek, & Banaji, 2003).

Explicit Measures

Using a 7-point scale ranging from 1 (*I strongly disagree*) to 7 (*I strongly agree*), we asked participants to indicate their agreement to two statements measuring how positive and how negative they evaluated group X ("I have a positive evaluation of group X"; "I have a negative evaluation of group X").

Rating Task

To quantify the extent to which classification images reflected meaningful variation caused by the behavioral information manipulation, 31 independent judges (four men; $M_{age} = 22.29$, $SD = 5.56$) rated the individual group X classification images on trustworthiness and criminality. They received €2 or course credits in return. In a block per trait, judges rated all images using a 9-point scale ranging from 1 (*not trustworthy/criminal*) to 9 (*very trustworthy/criminal*). The inter-rater reliability for both the trustworthiness and criminality ratings was high (Cronbach's alpha > .91). Trials in which judges responded faster than 300 milliseconds have been removed (<1%). We standardized ratings per participant. Subsequently, for each classification image, we calculated averaged trustworthiness and criminality ratings.

Procedure

Participants started with the stereotype formation task, in which they were asked to form an impression of group X and group Y based on group members' pictures and behavioral descriptions. It was explained to them that the behavioral descriptions represented typical behavior of the group member whose picture they were viewing and that random noise was added to pictures to make it more difficult to recognize the faces underneath.

[3]Note that the classification images can in principle be very different from the base images used. For instance, in Dotsch et al. (2008), a White Scandinavian base face was used, although the classification images clearly appeared Chinese or Moroccan, and in Dotsch et al. (2011), a male base face was used, although the resulting feminine classification image clearly appeared feminine.

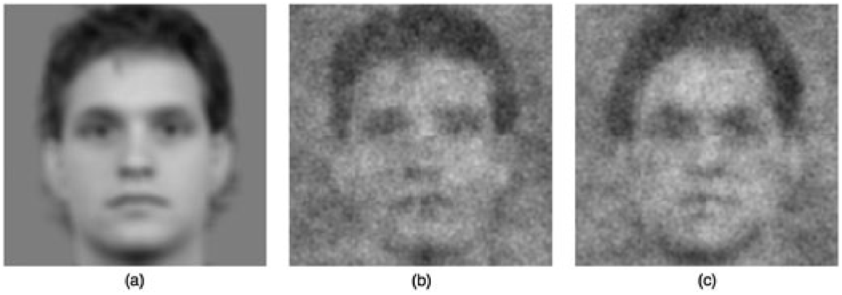

Figure 2. Averaged base face (a), example stimuli with noise superimposed (b) and inverse noise superimposed (c)

We aimed to create strong associations between groups and traits. Because transference of traits from individual group members to other members of the group is strongest for groups that are perceived to be highly entitative (e.g., Crawford et al., 2002), all participants received the same high entitativity instruction before each block of the stereotype formation task: "The members of group X [Y] are very similar to each other and do not differ in many ways from each other. The members come from similar backgrounds and have the same opinions, similar important beliefs, and similar personalities. Across a variety of situations, members of group X [Y] will act in a similar manner." (Crawford et al., p. 1080).

After the stereotype formation task, participants completed an unrelated filler task, which lasted on average about 18 minutes. Subsequently, they proceeded with the reverse correlation task, the IAT, and explicit measures. Afterwards, participants were debriefed and thanked.

Each participant's classification image was calculated, and, as explained previously, rated by independent judges on trustworthiness and criminality. These ratings were used as dependent variables in data analysis.

RESULTS

Data Preprocessing

Implicit Association Test

We processed the IAT data according to the D-measure algorithm (Greenwald et al., 2003). There were no participants with more than 10% of response latencies below 300 milliseconds. Trials with latencies above 10.000 milliseconds were excluded from analysis (0.04%). Latencies on trials with incorrect responses were replaced by the block mean latency of correct responses + 600 milliseconds. Latencies were then averaged for each block. We then computed difference scores for practice blocks (Block 4−Block 2) and target blocks (Block 5−Block 3), divided by their pooled SDs. The final D-score constituted the average of the two resulting difference scores. A positive D-score therefore indicated stronger positive than negative associations with group X (and stronger negative than positive associations with group Y).

Reverse Correlation Task

We generated classification images representing what participants expected typical group X members to look like (i.e., their visual group X representation) based on the reverse correlation data. Trials with response latencies below 300 milliseconds were excluded from analysis (2.20%). The classification images were calculated by averaging the sinusoid amplitude parameters of the selected noise patterns. The resulting average noise patterns were then superimposed on the base image to create the final classification image for each participant. In order to visually illustrate the effect of the behavioral information manipulation, we calculated the averaged classification images for the participants in the criminal and trustworthy conditions separately (Figure 3). The *individual* classification images were rated on criminality and trustworthiness by independent judges as described in the method section. For each image, we calculated the averaged criminality and trustworthiness ratings, which correlated $r = -.91$. We therefore combined the two ratings to compute a trustworthiness index by subtracting the averaged criminality ratings from the averaged trustworthiness ratings for each stimulus. A positive trustworthiness index therefore indicated that a classification images was rated as more trustworthy than criminal.

Exclusion Criteria

Two participants who scored more than 2.5 inter-quartile ranges below or above the median IAT D-score within their respective behavioral information conditions as well as across conditions were removed from analysis.[4] The final data set contained 75 participants.

Analyses

Analytical Strategy

For each of the dependent measures (classification image ratings, group X explicit evaluation, and group X implicit evaluation), we performed hierarchical linear regression where in a first step, the respective dependent measure was predicted by behavioral information condition (effects coded; 1 = trustworthy versus −1 = criminal) only. In a second step, to account for counter-balancing effects, we entered the base image used with group X faces (Base Image), noise set used with group X faces (Noise Set) and presentation order of the groups in the stereotype formation task (Order) and their first-order interactions with behavioral information condition. All counter-balancing variables were effects coded.

[4]When keeping all outliers in the data set all reported effects are still significant, with the exception of the effect of behavioral information condition on IAT D-scores.

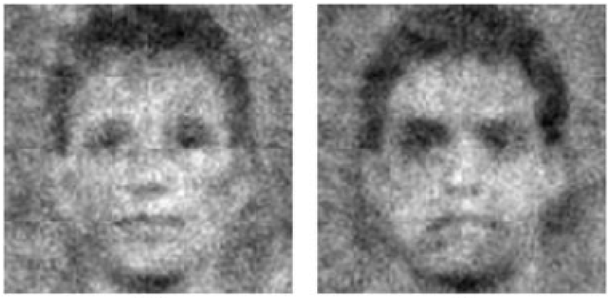

(a) (b)

Figure 3. Averaged classification images for trustworthy behavioral information condition (a) and for criminal behavioral information condition (b)

Classification Image Ratings

As described earlier, classification image ratings were standardized per judge prior to analysis, after which we computed a trustworthiness index for each stimulus by subtracting the averaged criminality ratings from the averaged trustworthiness ratings for that stimulus. As predicted, classification images of participants in the trustworthy behavioral information condition ($n = 37$) were rated as more trustworthy and less criminal ($M = 0.38$, $SD = 0.98$, $n = 38$) than classification images of participants in the criminal behavioral information condition ($M = -0.38$, $SD = 1.04$, $n = 37$), $\beta = 0.36$, $t(73) = 3.28$, $p < .01$, $R^2 = 0.13$. Entering the counter-balancing factors into the regression analysis did not increase R^2 significantly, $F(6, 67) < 1$, $p = .62$.

Group X Explicit Evaluation

Participants in the trustworthy behavioral information condition had more positive ($M_{\text{trustworthy}} = 6.08$, $SD = 1.19$ vs. $M_{\text{criminal}} = 2.95$, $SD = 1.88$) and less negative ($M_{\text{trustworthy}} = 2.79$, $SD = 2.06$ vs. $M_{\text{criminal}} = 5.14$, $SD = 1.90$) explicit evaluations of group X than those in the criminal behavioral information condition, respectively $\beta_{\text{positive}} = 0.71$, $t(73) = 8.62$, $p < .001$, $R^2 = 0.50$, and $\beta_{\text{negative}} = -0.51$, $t(73) = 5.13$, $p < .001$, $R^2 = 0.26$. Entering the counter-balancing factors into the regression analysis increased R^2 significantly for positive evaluation, to $R^2 = 0.64$, $F(6, 67) = 4.14$, $p < .01$, but not for negative evaluation, $F(6, 67) = 1.59$, $p = .16$. Importantly, for positive explicit evaluation, the effect of behavioral information condition was still significant in the more complex model, $\beta_{\text{positive}} = 0.54$, $t(67) = 5.93$, $p < .001$. All three counter-balancing variables interacted with the effect of behavioral information condition. Specifically, the behavioral information condition effect on positive explicit evaluation was stronger when group X was presented first, $\beta = .14$, $t(67) = 2.11$, $p = .04$; when the left base face in Figure 1(a) was used for group X, $\beta = .25$, $t(67) = 2.78$, $p < .01$; and weaker when Noise Set 1 was used for group X, $\beta = -.33$, $t(67) = 3.22$, $p < .01$. However, because the main effect of condition was significant in both the simple model and the model including the counter-balancing variables, and the interactions were only observed for positive explicit evaluation, the conclusion that behavioral information affected explicit evaluation remained.

Group X Implicit Evaluation

Participants' IAT D-scores also varied as a function of behavioral information condition, $\beta = -0.24$, $t(73) = 2.16$, $p < .03$, $R^2 = 0.06$. Participants in the trustworthy behavioral information condition had stronger positive than negative associations with group X ($M = 0.35$, $SD = 0.35$) compared with participants in the criminal behavioral information condition ($M = 0.16$, $SD = 0.41$). Entering the counter-balancing factors into the regression analysis did not increase R^2 significantly, $F(4, 67) < 1$, $p = .68$. However, the effect of behavioral information condition on IAT D-scores was no longer significant (and reversed sign), $\beta = 0.20$, $t(67) = 1.40$, $p = .17$. This was presumably the joint result of an initially small effect and the loss of degrees of freedom due to the additional predictors.

Correlations between Measures

See Table 1 for a correlation matrix of all variables. Although positive and negative explicit evaluations correlated with participants' implicit evaluation, surprisingly, ratings of classification image trustworthiness and criminality (trustworthiness index) correlated only with participants' explicit evaluation. That is, participants who evaluated group X more positively or less negatively on an explicit level created classification images that were rated as more trustworthy and less criminal. This was not observed on the implicit level.

DISCUSSION

Using a stereotype formation task, we tested whether verbal behavioral information diagnostic of an out-group's traits biases perceivers' visual expectations of typical out-group faces toward having facial features corresponding with the inferred traits. Our hypothesis was confirmed: classification

Table 1. Correlation matrix of all variables

	1	2	3	4
1. Manipulation (1 = trustworthy, −1 = criminal)				
2. Classification image trustworthiness index	**0.36**			
3. Positive explicit evaluation	**0.71**	**0.46**		
4. Negative explicit evaluation	**−0.51**	**−0.38**	**−0.73**	
5. Implicit evaluation (IAT D-score, higher is more positive implicit evaluation)	**0.24**	.08	**.32**	**−.29**

Note: IAT, Implicit Association Test.
Correlations in bold are $p < .05$.

images of participants in the criminal behavioral information condition were judged to be more criminal than those of participants in the trustworthy behavioral information condition. Because the presented exemplar faces in the formation task were kept constant (and counter-balanced), the reported bias in the expected facial appearance of group members could be caused only by exemplars' behavioral information. Note that participants' expectations for typical group X faces could simply have resembled the base face used for group X in the formation task. Instead, behavioral information about the exemplars influenced the expected facial appearance of a typical group X member. To our knowledge, the current work is the first in demonstrating experimental effects of behavioral information on the expected facial appearance of out-group members. Previous work has made use of existing groups, because of which the bias in expected facial appearance could have been a result of the actual facial appearance of previously encountered group members. Here, we demonstrated that behavioral information unrelated to facial appearance is sufficient to generate this bias.

The current work highlights the usefulness of reverse correlation tasks to assess spontaneous face processing biases. However, the outcome of the reverse correlation task, a classification image, should not be equated with visual stereotype content, however tempting. In fact, here, the classification image is interpreted as what participants expected faces of typical group members to look like. These expectations might have been drawn either from a mental visual image or from verbal knowledge about the group. Below, we briefly expand on these two options.

In the first account, the mental visual image account, it is assumed that participants, on the basis of verbal information, imagined what typical out-group faces looked like and encoded a biased visual image into their stereotype. Thus, parallel to integrating the verbal behavioral information into a more abstract verbal stereotype (i.e., either criminal or trustworthy), participants may also have formed a visual mental image of the group's typical facial appearance partly on the basis of the exemplar faces provided. The crucial assumption is that in the process of forming a visual image, the concurrent processing of encoding and integrating behavioral information may already have biased the mental representation of the group's facial features in line with the stereotype. In the subsequent reverse correlation task, the participants could then simply compare the presented stimulus faces with their mental image of the group's typical face and choose the one with the most feature overlap.

In the alternative account, the verbal stereotype account, it is assumed that participants' decisions in the reverse correlation task are merely based on the abstract verbal stereotype they have formed of the group, that is, either criminal or trustworthy. In choosing between two stimulus faces, participants choose the one that looks more criminal or trustworthy. The verbal stereotype then becomes a heuristic tool to choose between the two stimulus faces without having to resort to a mental visual image of group X. However, in this case, participants use a mental visual image of what a criminal or trustworthy face looks like.

On the basis of our present data, we cannot decide which of these two accounts is the most plausible one. However, we slightly favor the mental visual image account, for various reasons. First and most importantly, while the verbal stereotype account may seem the most parsimonious one at first sight (because it circumvents the assumption of the formation of a visual representation of group X), it is actually more cumbersome by positing a more complex decision process. That is, in the verbal stereotype account, deciding between two stimulus faces requires that participants first retrieve the verbal stereotype, then its concomitant visual stereotype and finally use the latter as a heuristic device to decide which of two faces most likely is from group X. In contrast, the mental visual image account only requires the additional assumption of the parallel formation of a visual representation of group X to make for a simpler direct facial feature matching decision process. Second, in real life people can identify virtually immediately members of specific categories in a split second. It seems implausible that they rely on verbal stereotype activation and subsequent retrieval of the facial features belonging to the stereotype to accomplish this fast social category identification. The most plausible account of this identification process seems to be that people have mental visual representations of groups in which features associated with their stereotypical expectations have been incorporated, probably because in the formation of these visual stereotypes, facial features without specific stereotypic connotations and trait-related features have merged into a holistic representation face representation of the group in question (see Dotsch et al., 2011). It is quite conceivable that, in an embryonic form, the participants in the present study similarly formed a visual image of group X in which features of the exemplar faces and features of the induced trait have similarly merged into a single face representation. Considering these arguments, we slightly favor the mental visual image account over the verbal stereotype account. However, as indicated earlier, our data do not provide conclusive evidence concerning either explanation.

Participants' explicit and implicit attitudes varied in line with the behavioral information manipulation. This replicates work by Ratliff and Nosek (2010), who induced both implicit and explicit attitudes in a similar stereotype formation paradigm, using personality traits instead of behavioral information, but without the assessment of effects on the expected facial appearance of group members. The behavioral information manipulation had a stronger effect on explicit attitudes than on implicit attitudes. This fits theories stating that the implicit system learns more slowly than the explicit system (Rydell & McConnell, 2006; Sloman, 1996; Smith & DeCoster, 2000). Perhaps leaving more time between stereotype formation and attitude assessment or presenting more group exemplars would increase the effect of behavioral information on implicit attitudes. Additionally, the IAT might be a noisier measure than explicit ratings. For instance, the IAT score not only reflects implicit attitudes but also individual differences in task switching capacity (Klauer, 2005; Mierke & Klauer, 2001), generating spread across participants unrelated to the behavioral information manipulation.

The IAT D-score showed a positivity bias (i.e., the average IAT D-score was not negative in the criminal behavioral information condition). One possible source of this bias is the fixed IAT block order that we used to minimize between-participant variation. Group X was always paired first with positive stimuli (and Y with negative). This may have slightly strengthened the association of group X with positive valence (and Y with negative) as a result of which the later reversal (i.e., pairing group X with negative and Y with positive) may have yielded slower responses. Thus, this order of IAT blocks may have unintentionally produced slightly more positive evaluations of group X and slightly more negative evaluations of group Y. Another source of bias might be the fact that the IAT, pitting group X against group Y, was administered after participants performed the extensive reverse correlation task, which was focused only on group X. This greater exposure to group X than to group Y might have induced a positivity bias toward group X in a way similar to the implicit partisanship effect (Greenwald, Pickrell, & Farnham, 2002; Pinter & Greenwald, 2004). A third reason might be the fact that we had a predominantly female sample, and the faces in the stereotype formation phase were all generated using a male base face. This might have biased the implicit evaluation of group X in a general direction. Nonetheless, and more importantly, people had relatively more negative implicit attitudes in the criminal behavioral information condition than in the trustworthy behavioral information condition.

Unexpectedly, implicit attitudes did not predict the expected out-group facial appearance. This means we did not conceptually replicate Dotsch et al. (2008), in which more implicitly prejudiced participants (i.e., participants with stronger negative implicit attitudes) demonstrated more criminal-looking mental representations of typical out-group faces. However, explicit attitudes did correlate with the expected out-group facial appearance. Dotsch et al. (2008) did not report any correlations between explicit attitudes and visualized representations of Moroccan faces. Re-analysis of their data (Dotsch et al., Study 2) for the purpose of the current work showed that neither positive nor negative explicit attitudes correlated with the criminality or trustworthiness judgments of participants' visualized representations of Moroccan faces. This might be explained by a social norm not to appear explicitly prejudiced toward real-world groups (Crandall, Eshleman, & O'Brien, 2002; Dovidio & Gaertner, 2004). This norm might not be in place for the novel and neutrally labeled group X in the current study.

The absence of the correlation between implicit attitudes and expected out-group facial appearance might be explained by the slow-learning mechanism mentioned previously. The learning history for the new group might simply not (yet) be comparable with the extensive and culturally embedded learning history for existing groups, such as Moroccans. Another potential explanation might be the unusual character of the present IAT, in which the stimuli for the group X and group Y categories were the group X and group Y labels themselves. Generally, IATs use a more varied set of words that are representative of the categories of interest. Furthermore, the Dotsch et al. study used a single target IAT (Bluemke & Friese, 2008; Karpinski & Steinman, 2006) as a measure of implicit association strength, in which no reference category is used, instead of a traditional two-category IAT. This methodological decision might have contributed to the absence of the previously observed correlation.

In the stereotype formation task, each group was assigned only one trait, opposite in valence. Likewise, our judges were instructed to rate classification images on the same valenced traits. It is therefore unclear whether the classification images and ratings reflected mere valence or the actual traits criminal or trustworthy. Importantly, we acknowledge the equivocal specificity of our results (valence versus specific traits) but emphasize that the mechanism (that behavioral information affected the expected facial appearance of group members) remains the same, regardless of the effects taking place on the level of valence or traits. Moreover, note that previous work on reverse correlation technique successfully visualized representations of John Travolta and Tom Cruise (Mangini & Biederman, 2004), Mona Lisa (Kontsevich & Tyler, 2004), or the orthogonal dimensions of trustworthiness (valence) and dominance (Dotsch & Todorov, 2011), demonstrating that classification images are not strictly limited to conveying valence. Future research might tease apart the influence of valence and specific traits (similar to Amodio & Devine, 2006; Dotsch et al., 2011), for instance, by inducing multidimensional stereotypes which vary on more dimensions than just the valence dimension. Furthermore, it might not even be necessary to associate groups with specific stereotype content as classifying faces as merely an out-group might already be sufficient for a negativity bias to emerge in face perception. Work on this intriguing possibility has already been initiated (Ratner, Dotsch, Wigboldus, van Knippenberg, & Amodio, in preparation).

Here, because of our programmatic interest in how group stereotypes influence face perception, we focused on the influence of behavioral information on the expected facial representation of *group members*. However, on first sight, it seems probable that the same processes are at work in the expectations that people form about *individuals*. It has been shown that knowledge about a person's personality shapes the perception of facial features and perceived similarity of faces (Hassin & Trope, 2000). The current work builds on these earlier findings by directly tapping into the effects of behavioral information on visual representations. However, we also extend the earlier findings to the group level, which is not trivial. Forming expectations about the facial appearance of group members might be a qualitatively different process than forming expectations about the facial appearance of individuals. For group-based expectations, information about multiple individuals needs to be integrated into a single stereotype, whereas for individual-based expectations no such information integration process is required. As a consequence, any findings about forming impressions about individuals do not directly generalize to what happens during forming stereotypes about groups.

The major contribution of the present study is that it establishes a causal link between behavioral information about exemplars and the expected facial appearance of group members. It demonstrates that verbal information unrelated to visual appearance can affect visual representations of groups in a spontaneous manner. When people witness members of a specific novel group commit a criminal act, what they learn is not only that the group might behave criminally but also that group members' faces look criminal, even though the original faces might not have looked criminal at all.

REFERENCES

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* (4), 652–661.

Barsalou, L. W. (1999). Perceptual symbol systems. *The Behavioral and Brain Sciences*, *22*(4), 577–609.

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.

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

Brewer, M. B. (1988). A dual process model of impression formation. In T. K. Srull, & R. S. Wyer (Eds.), *Advances in social cognition* (Vol. 1, pp. 1–36). Hillsdale, NJ: Erlbaum.

Cacioppo, J. T., & Petty, R. E. (1982). The need for cognition. *Journal of Personality and Social Psychology*, *42*, 116–131.

Campbell, D. T. (1958). Common fate, similarity, and other indices of the status of aggregates of persons as social entities. *Behavioral Science*, *3*, 14–25.

Carlston, D. E. (1994). Associated systems theory: A systematic approach to cognitive representations of persons. In T. K. Srull, & R. S. Wyer (Eds.), *Advances in social cognition* (Vol. *7*, pp. 1–78). Hillsdale, NJ: Erlbaum.

Castelli, L., Carraro, L., Gawronski, B., & Gava, K. (2010). On the determinants of implicit evaluations: When the present weighs more than the past. *Journal of Experimental Social Psychology*, *46*(1), 186–191.

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.

Crandall, C. S., Eshleman, A., & O'Brien, L. (2002). Social norms and the expression and suppression of prejudice: The struggle for internalization. *Journal of Personality and Social Psychology*, *83*(3), 359–378.

Cratylus. (1995). *Need for closure (Dutch version)*. Amsterdam: Vrije Universiteit: Vakgroep Sociale Psychologie.

Crawford, M. T., Sherman, S. J., & Hamilton, D. L. (2002). Perceived entitativity, stereotype formation, and the interchangeability of group members. *Journal of Personality and Social Psychology*, *83*(5), 1076–1094.

De Houwer, J., Teige-Mocigemba, S., Spruyt, A., & Moors, A. (2009). Implicit measures: A normative analysis and review. *Psychological Bulletin*, *135*(3), 347–368.

Dijksterhuis, A., & van Knippenberg, A. (1999). On the parameters of associative strength: Central tendency and variability as determinants of stereotype accessibility. *Personality and Social Psychology Bulletin*, *25* (4), 529–536.

Dotsch, R., & Todorov, A. (2011). Reverse correlating social face perception. *Social Psychological and Personality Science*. Downloaded online on 12 December 2011 from http://spp.sagepub.com/content/early/2011/12/08/1948550611430272.

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.

Dotsch, R., Wigboldus, D. H. J., & van Knippenberg, A. (2011). Biased allocation of faces to social categories. *Journal of Personality and Social Psychology*, *100*, 999–1014.

Dovidio, J. F., & Gaertner, S. L. (2004). Aversive racism. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 36, pp. 1–52). San Diego, CA: Academic Press.

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

Ford, T. E., & Stangor, C. (1992). The role of diagnosticity in stereotype formation: Perceiving group means and variances. *Journal of Personality and Social Psychology*, *63*, 356–367.

Freeman, J. B., & Ambady, N. (2011). A dynamic interactive theory of person construal. *Psychological Review*, *118*(2), 247–279.

Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, *132*(5), 692.

Gluszek, A., & Dovidio, J. F. (2010). The way they speak: A social psychological perspective on the stigma of nonnative accents in communication. *Personality and Social Psychology Review*, *14*(2), 214–237.

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.

Gosselin, F., & Schyns, P. G. (2003). Superstitious perceptions reveal properties of internal representations. *Psychological Science*, *15*(5), 505–509.

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

Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology*, *85*, 197–216. doi: 10.1037/0022-3514.85.2.197

Greenwald, A. G., Pickrell, J. E., & Farnham, S. D. (2002). Implicit partisanship: Taking sides for no reason. *Journal of Personality and Social Psychology*, *83*, 367–379.

Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., & Banaji, M. R. (2009). Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity. *Journal of Personality and Social Psychology*, *97* (1), 17–41.

Hamilton, D. L., & Gifford, R. K. (1976). Illusory correlation in intergroup perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology*, *12*, 392–407.

Hamilton, D. L., & Sherman, S. J. (1989). Illusory correlations: Implications for stereotype theory and research. In D. Bar-Tal, C. F. Graumann, A. W. Kruglanski, & W. Stroebe (Eds.), *Stereotyping and prejudice. Changing conceptions* (pp. 59–82). New York: Springer-Verlag.

Hamilton, D. L., & Sherman, S. J. (1996). Perceiving persons and groups. *Psychological Review*, *103*, 336–355.

Hassin, R., & Trope, Y. (2000). Facing faces: Studies on the cognitive aspects of physiognomy. *Journal of Personality and Social Psychology*, *78*(5), 837–852.

Higgins, E. T., & Bargh, J. A. (1987). Social cognition and social perception. *Annual Review of Psychology*, *38*, 369–427.

Hirschfeld, L. A. (1995). Do children have a theory of race? *Cognition*, *54*, 209–252.

Hugenberg, K. (2005). Social categorization and the perception of facial threat: Target race moderates the response latency advantage for happy faces. *Emotion*, *5*, 267–276.

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.

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.

Imhoff, R., Dotsch, R., Bianchi, M., Banse, R., & Wigboldus, D. H. J. (2011). Facing Europe: Visualizing spontaneous ingroup projection. *Psychological Science*, *22*, 1583–1590.

Jussim, L. (1991). Social perception and social reality: A reflection-construction model. *Psychological Review*, *98*, 54–73.

Karpinski, A., & Steinman, R. B. (2006). The single category Implicit Association Test as a measure of implicit social cognition. *Journal of Personality and Social Psychology*, *91*, 16–32.

Klauer, K. C. (2005). Task-set inertia, attitude accessibility, and compatibility-order effects: New evidence for a task-switching account of the implicit association test effect. *Personality and Social Psychology Bulletin*, *31*(2), 208–217.

Kontsevich, L., & Tyler, C. W. (2004). What makes Mona Lisa smile. *Vision Research*, *44*(13), 1493–1498. doi: 10.1016/j.visres.2003.11.027

Langner, O., Dotsch, R., Bijlstra, G., Wigboldus, D. H. J., Hawk, S., & Van Knippenberg, A. (2010). Presentation and validation of the Radboud Face Database. *Cognition & Emotion*, *24*(8), 1377–1388.

Lewicki P. (1986). *Nonconscious social information processing*. Orlando, FL: Academic Press.

Lippmann, W. (1922). *Public opinion*. New York, NY: Harcourt Brace.

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

McArthur, L. Z., & Baron, R. M. (1983). Toward an ecological theory of social perception. *Psychological Review*, *90*, 215–238.

McGarty, C. (2002). Stereotype formation as category formation. In C. McGarty, V. Y. Yzerbyt, & R. Spears (Eds.), *Stereotypes as explanations: The formation of meaningful beliefs about social groups* (pp. 16–37). University Press: Cambridge.

Mierke, J., & Klauer, K. C. (2001). Implicit association measurement with the IAT: Evidence for effects of executive control processes. *Zeitschrift für Experimentelle Psychologie*, *48*(2), 107–122.

Monteith, M. J., Voils, C. I., & Ashburn-Nardo, L. (2001). Taking a look underground: Detecting, interpreting, and reacting to implicit racial bias. *Social Cognition*, *19*, 395–417.

Mullen, B., & Johnson, C. (1990). Distinctiveness-based illusory correlations and stereotyping: A meta-analytic integration. *British Journal of Social Psychology*, *29*, 11–28.

Pinter, B., & Greenwald, A. G. (2004). Exploring implicit partisanship: Enigmatic (but genuine) group identification and attraction. *Group Processes & Intergroup Relations*, *7*(3), 283–296.

Ratliff, K. A., & Nosek, B. A. (2010). Creating distinct implicit and explicit attitudes with an illusory correlation paradigm. *Journal of Experimental Social Psychology*. doi: 10.1016/j.jesp.2010.04.011

Ratner, K. G., Dotsch, R., Wigboldus, D. H. J., van Knippenberg, A., & Amodio, D. M. (in preparation). Ingroup favoritism in the mind's eye: Visual representations of faces following minimal group assignment. In preparation.

Rudman, L. A. (2004). Sources of implicit attitudes. *Current Directions in Psychological Science*, *13*(2), 79–82.

Rudman, L. A., Phelan, J. E., & Heppen, J. B. (2007). Developmental sources of implicit attitudes. *Personality and Social Psychology Bulletin*, *33*, 1700–1713.

Rydell, R. J., & McConnell, A. R. (2006). Understanding implicit and explicit attitude change: A systems of reasoning analysis. *Journal of Personality and Social Psychology*, *91*(6), 995–1008.

Sherman, J. W. (1996). Development and mental representation of stereotypes. *Journal of Personality and Social Psychology*, *70*, 1126–1141.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, *119*, 3–22.

Smith, E. R., & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review*, *4*, 108–131.

Smith, E. R., & Zarate, M. A. (1990). Exemplar and prototype use in social categorization. *Social Cognition*, *8*, 243–262.

Todorov, A., Dotsch, R., Wigboldus, D. H. J., & Said, C. P. (2011). Data-driven methods for modeling social perception. *Social and Personality Psychology Compass*, *5*(10), 775–791. doi: 10.1111/j.1751-9004.2011.00389.x

Todorov, A., Said, C. P., Engell, A., & Oosterhof, N. N. (2008). Understanding evaluation of faces on social dimensions. *Trends in Cognitive Science*, *12*, 455–460.

Wigboldus, D. H. J., Semin, G. R., & Spears, R. (2000). How do we communicate stereotypes? Linguistic bases and inferential consequences. *Journal of Personality and Social Psychology*, *78*(1), 5–18.

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: Guilford.

Zebrowitz, L. A., & Montepare, J. M. (2008). Social psychological face perception: Why appearance matters. *Social and Personality Psychology Compass*, *2*(3), 1497–1517.

