Published in final edited form as:

J Pers Soc Psychol. 2010 February ; 98(2): 175–189. doi:10.1037/a0017990.

Facial Resemblance to Emotions: Group Differences, Impression Effects, and Race Stereotypes

Leslie A. Zebrowitz,

Department of Psychology, Brandeis University

Masako Kikuchi, and

Department of Psychology, Brandeis University

Jean-Marc Fellous

Department of Psychology, University of Arizona

Abstract

The authors used connectionist modeling to extend previous research on emotion overgeneralization effects. Study 1 demonstrated that neutral expression male faces objectively resemble angry expressions more than female faces do, female faces objectively resemble surprise expressions more than male faces do, White faces objectively resemble angry expressions more than Black or Korean faces do, and Black faces objectively resemble happy and surprise expressions more than White faces do. Study 2 demonstrated that objective resemblance to emotion expressions influences trait impressions even when statistically controlling possible confounding influences of attractiveness and babyfacedness. It further demonstrated that emotion overgeneralization is moderated by face race and that racial differences in emotion resemblance contribute to White perceivers' stereotypes of Blacks and Asians. These results suggest that intergroup relations may be strained not only by cultural stereotypes but also by adaptive responses to emotion expressions that are overgeneralized to groups whose faces subtly resemble particular emotions.

Keywords

face perception; emotion overgeneralization; impression formation; race stereotypes

Let's not forget that the little emotions are the great captains of our lives and we obey them without realizing it.

—Van Gogh, 1889

Although Van Gogh was referring to obeying one's own emotions, the fact is that people are also influenced by others' "little" emotions without realizing it. Impressions of people's stable traits are influenced not only by their full-blown emotional expressions but also by the subtle resemblance of their permanent facial structure to an emotional expression. The first of these effects has been identified as *temporal extension*: "the perceiver regards a momentary characteristic of the person as if it were an enduring attribute" (Secord, 1958, p. 313). Consistent with this hypothesis, people who display transient angry expressions are perceived to have stable traits associated with low warmth and high dominance; those with sad, fearful, or surprised expressions are perceived as moderate in warmth and low in

dominance; and those with happy expressions are perceived as high in warmth and dominance (Hess, Blairy, & Kleck, 2000; Knutson, 1996; Montepare & Dobish, 2003; Zebrowitz, Kikuchi, & Fellous, 2007). The second effect, whereby trait impressions are influenced by the resemblance of a person's permanent facial structure to an emotional expression, was identified by Zebrowitz as *emotion overgeneralization* (Zebrowitz, 1996, 1997). This phenomenon raises the question of what specific kinds of neutral expression faces resemble an emotion expression. Previous research suggests that faces varying in sex or facial maturity differ in resemblance to emotions. In the present article, we present the results of two studies in which connectionist modeling is used to extend previous research on such emotion overgeneralization effects. In Study 1, we assessed similarities between emotion expressions and neutral expression faces varying in race and sex¹; in Study 2, we determined whether trait impressions of individuals and race stereotypes are influenced by variations in resemblance to emotion expressions.

Study 1: Effects of Face Sex and Race on Neutral Faces' Resemblance to Emotion Expressions

A number of studies have demonstrated that faces from certain demographic categories differ in their resemblance to particular emotion expressions. Faces of babies resemble surprise and fear more than do faces of adults, and baby faces show less resemblance to anger, effects that have been demonstrated both by subjective ratings of the faces (Marsh, Adams, & Kleck, 2005) and by connectionist modeling using facial metrics (Zebrowitz, Kikuchi, & Fellous, 2007). In addition, fear was more accurately recognized and anger less accurately recognized in adult faces manipulated to have a more babyish structure—larger eyes or rounder face (Sacco & Hugenberg, 2009). Consistent with evidence that women's faces are more neotenous than men's (Enlow, 1990), surprise and angry expressions are also differentially associated with female and male faces, with surprised faces of either sex judged more feminine than angry ones, and the emotion of female surprised faces identified more quickly than that of female angry faces, with the reverse for male faces (Le Gal & Bruce, 2002). Similar results are shown in research assessing the resemblance of female and male faces to angry versus happy expressions (Becker, Kenrick, Neuberg, Blackwell, & Smith, 2007; Hess, Blairy, & Kleck, 1997). Perceivers showed faster recognition of angry expressions in male than in female faces and more accurate and faster recognition of angry than happy faces as male, whereas they showed faster recognition of happy expressions in female than in male faces and faster recognition of happy than angry faces as female. In addition, neutral expression male faces were more often misidentified as angry than were neutral expression female faces, and androgynous faces morphed in a masculine direction were judged as angrier than those morphed in a feminine direction even though the expressions were identical.

There are two possible explanations for overlapping responses to different emotion expressions and to faces that vary in sex. One is that there are similarities between the *cultural meaning* of the demographic categories and the emotion expressions. Indeed, people rate men as more likely to display anger (Hess et al., 1997). Another possible explanation is that there are *structural similarities* between faces from certain demographic categories and those with certain emotion expressions. Clear evidence for an influence of facial structure on emotion identification is provided by the finding that the same faces are perceived as angrier when they are manipulated to have a shorter vertical distance between the eyes and mouth

¹It should be acknowledged that anthropologists and biologists question the validity of race as a scientific concept (e.g., Lewontin, 2001). Nevertheless, race is a widely accepted concept in folk psychology (Zuckerman, 1990). Although recognizing that clear decisions on category membership are problematic, we nevertheless use the "fuzzy" category system of racial groups and the terms *White*, *Korean*, and *Black* to denote the physical appearance of target faces.

(Neth & Martinez, 2009). More directly relevant to demographic differences in resemblance to emotion expressions, Marsh et al. (2005) argued that the morphology of emotion expressions derives from the adaptive utility of their mimicking variations in facial maturity. More specifically, they argued that fear and anger expressions evolved to mimic baby's faces and mature faces, respectively. Their rationale was that it is adaptive for those experiencing fear to elicit reactions paralleling those elicited by helpless babies and for those experiencing anger to elicit reactions paralleling those elicited by powerful adults. The fact that a connectionist model detected similarities between faces of babies and surprise expressions and faces of adults and angry expressions (Zebrowitz, Kikuchi, Fellous, 2007) supports a structural similarities explanation for overlapping responses to surprise versus anger and babies versus adults, because the model has no knowledge of overlapping social stereotypes about people of certain ages and emotional states. Because adult female faces retain more neotenous qualities than do male faces, these results also favor a *structural similarities* explanation for sex differences in facial resemblance to surprise and anger, which should parallel the documented age differences. Indeed, Becker et al. (2007) proposed that angry expressions evolved to mimic the signal of maleness that is communicated by a heavier brow.

In contrast to research investigating the intersection of different emotion expressions with faces differing in age or sex, research investigating overlapping responses to particular emotion expressions and faces varying in race has focused on similarities in the cultural meaning of the racial categories and the emotion expressions as an explanatory mechanism. One study found that European Americans high in racial prejudice were faster to recognize the onset of anger and slower to recognize the offset of anger in Black faces than in White faces, whereas this was not true for those low in prejudice (Hugenberg & Bodenhausen, 2003). However, this study is not well equipped to reveal any commonalities between race-related facial structure and emotion expressions, because the Black and White schematic faces were identical except for their skin tone. Another study, using computer-generated racially ambiguous faces, found that European Americans high in racial prejudice were more likely to categorize angry than happy faces as African American, whereas this was not true for those low in prejudice (Hugenberg & Bodenhausen, 2004). Again, this study is better suited to capturing the overlapping cultural meaning of emotion expressions and racial category rather than to elucidating any overlap in the racial and expression morphology, because the racial morphology was ambiguous by design. However, there are racial differences in facial structure that could yield differences in resemblance to emotion expressions. For example, Asian faces have narrower eyes than White or Black faces, and, as noted above, fear was less accurately recognized and anger more accurately recognized in adult faces manipulated to have smaller eyes (Sacco & Hugenberg, 2009). However, Asian faces were judged more baby-faced than White faces (Zebrowitz, Bronstad, & Lee, 2007), perhaps because they have rounder rather than longer faces, a feature that leads to more accurate recognition of fear and less accurate recognition of anger (Sacco & Hugenberg, 2009).

We used connectionist modeling in Study 1 to determine whether there are structural similarities between certain demographic categories and certain emotion expressions. The use of connectionist models has an advantage over previous research methods because such models are impervious to stereotyped assumptions. For example, if a network trained to differentiate angry expressions from neutral expressions subsequently finds a new set of male or Black faces more physically similar to the anger expressions than female or White faces, then this cannot be attributed to knowledge of overlapping social stereotypes about angry people and men or African Americans. Rather, it can be due only to intrinsic similarities in the facial attributes known to the neural network. It should be noted that connectionist modeling is used in this research as a mathematical technique for generating

an objective index of the structural similarity of a face to a particular category of faces. The aim of this research was not to test various alternative cognitive or neural models of face processing.

If structural similarities between male or female faces and certain emotion expressions can account for previous evidence that face sex moderates responses to emotion expressions and vice versa, then we predict the following:

- 1 Compared with male faces, female faces will elicit greater activation of neural network output units trained to detect surprise or happy expressions.
- 2 Compared with female faces, male faces will elicit greater activation of a neural network output unit trained to detect angry expressions.

However, if similarities in the *cultural meaning* of male or female faces and certain emotion expressions account for previously documented moderating effects, then neural networks trained to detect emotion expressions should not respond differentially to male and female faces.

- 3 Because the existing research on face race and emotion expressions was designed to examine the influence of cultural stereotypes rather than facial structure, our investigation of effects of face race on the activation of output units trained to respond to surprised, happy, or angry faces was exploratory rather than guided by clear predictions.

Method

Faces

Connectionist models were trained to recognize the facial metrics of angry, happy, and surprise expressions in White male and female training/test faces. The extent to which the models detected similarities to these emotions in neutral expression faces was then examined using a separate set of White, Black, and Asian generalization faces.

Training/test faces—Training/test faces were taken from a previous study (Zebrowitz, Kikuchi, & Fellous, 2007). They included digitized black-and-white portrait photos of 26 Caucasian men and 26 Caucasian women in their 20s or 30s, each of whom posed neutral, happy, angry, and surprise expressions. Previously reported validations of the database demonstrated significantly higher ratings of anger for angry faces than each of the other categories, higher ratings of happy for happy faces than each of the other categories, and higher ratings of surprise for surprise faces than each of the other categories (Zebrowitz, Kikuchi, & Fellous, 2007).

Generalization faces—Generalization faces included 120 White, 120 Korean, and 120 Black neutral expression faces, with men and women equally represented within each race. White facial images were selected from four different databases: University of Stirling PICS database (<http://pics.psych.stir.ac.uk/>); the AR face database (Martinez & Benavente, 1998); NimStim²; and two year-books (one high school and one university yearbook). Korean facial images were selected randomly from a Korean university yearbook. The facial images of Black men were selected from a set of faces that had been used in a study by Blair, Judd, Sadler, and Jenkins (2002). The images of the Black female faces were selected from the

²Development of the MacBrain Face Stimulus Set was overseen by Nim Tottenham and supported by the John D. and Catherine T. MacArthur Foundation Research Network on Early Experience and Brain Development. Please contact Nim Tottenham at tott0006@tc.umn.edu for more information concerning the Stimulus Set.

website <http://americansingles.com> by searching for Black females ages 18–25. Four criteria were used for image selection: neutral expression, no head tilt, no eyeglasses, and no beards. Images were edited in Photoshop to remove jewelry and blemishes. All faces were shown in color against an identical beige background. To verify that faces had neutral expressions, four White judges (2 men) provided smile ratings on a 5-point scale, with endpoints labeled no smile/big smile. Faces were blocked by sex and presented randomly within each block.

Facial metrics

Following the procedure reported by Zebrowitz, Kikuchi, & Fellous (2007), in-house software was used to mark 64 points on digitized images of each face from which facial metrics were computed using automatic procedures written in Visual Basic and Excel (see Figure 1). After establishing reliability for points marked by two judges on a subset of each category, one judge marked the remaining faces, and those points were used to calculate the final facial metrics. Eighteen nonredundant facial metrics that achieved acceptable reliability were selected as full facial inputs to the connectionist model. These included facial roundness plus the metrics shown in Figure 2. Facial roundness was computed by determining the average of the radii of two circles—one created by connecting facial points 31 right, 35 right, and 12, and another circle connecting facial points 31 left, 35 left, and 12, with a smaller average radius signifying more roundness. To adjust for variations in distance from the camera, each facial metric was normalized by interpupil distance (E2).

Connectionist modeling

The *total set* of faces used to train the network was composed of 208 faces (52 each for neutral, surprise, angry, and happy expressions). Networks were trained to differentiate either happy from neutral faces, surprise from neutral faces, or angry from neutral faces. On each of 20 trials, 34 neutral and 34 emotion faces (either happy, surprise, or angry faces) were randomly selected from the total set to compose the *training set*, with a different random set of faces selected on each trial. The rest composed the *test set*. The modeling had three phases. First, in the *training phase*, the 18 reliable facial metrics were provided as input to artificial neural networks that were trained with supervised learning to differentiate the 68 training faces (34 men), half with a neutral expression and half with an emotion expression. In the second or *test phase*, the trained network was tested on the set of 36 test faces (18 men) that differed in the emotion on which the network had been trained in order to establish that training was successful. In the third phase, the *generalization phase*, the trained network was provided with input metrics from the 360 neutral expression generalization faces, and the extent to which the output units responded to each of these faces was determined. These three phases were repeated for 20 trials to establish a reliable index of network activation by each face. Performing the entire procedure for networks trained to differentiate neutral expression faces from each of the three different emotion expressions generated three dependent variables for each generalization face: average activation across 20 trials of the happy output unit, the surprise output unit, and the angry output unit.

The connectionist models were standard back-propagation neural networks with one input layer, one hidden layer, and one output layer. Each input node projected to any or all of the hidden nodes, and the hidden nodes projected to the two output units (neutral and one of the emotions). The input weight matrices connecting the layers consisted of numbers between –1 and 1. The output units were rescaled into graded values ranging from 0% to 100% activation. All units were nonlinear and mapped the weighted sum of their inputs to their output using a sigmoidal transfer function. The training parameters were four hidden nodes, 3,000 training epochs per trial, a .02 learning rate, and a .2 error goal.

Results

Reliability of facial metrics

High interjudge agreement for the facial metrics of the emotion expression faces (cf. training and test set) was previously reported by Zebrowitz, Kikuchi, and Fellous (2007). The selected input metrics for the faces in the generalization set also showed high agreement both for the normalization interpupillary distance ($r = .97$) and the selected input metrics (mean $r = .85$).

Network training

The various networks are discussed below.

Happy-neutral networks—Training a network to differentiate happy and neutral faces achieved 94.56% correct identification of the 68 training faces and 86.81% correct identification of the 36 test faces, averaged across 20 trials, with activation of the happy unit significantly higher for happy faces ($M = 84.60$, $SD = 18.14$) than for neutral ones ($M = 13.61$, $SD = 11.65$), $F(1, 102) = 563.82$, $p < .0001$.

Surprise-neutral networks—Training a network to differentiate surprise and neutral faces achieved 94.41% correct identification of the 68 training faces and 87.92% correct identification of the 36 test faces, averaged across 20 trials, with significantly higher activation of the surprise unit by surprise faces ($M = 86.65$, $SD = 15.76$) than neutral ones ($M = 14.83$, $SD = 13.48$), $F(1, 102) = 623.55$, $p < .0001$.

Anger-neutral networks—Training a network to differentiate angry and neutral faces achieved 89.26% correct identification of the 68 training faces and 68.06% correct identification of the 36 test faces, averaged across 20 trials, with activation of the angry unit significantly higher for angry faces ($M = 72.99$, $SD = 26.96$) than for neutral ones ($M = 21.02$, $SD = 16.93$), $F(1, 102) = 138.56$, $p < .0001$. It should be noted that the less successful training of the anger- than the happy- or surprise-neutral networks is consistent with human judges' ratings of the faces. Zebrowitz, Kikuchi, and Fellous (2007) found that the neutral faces were rated higher in anger than in happiness or surprise. Also, although they were rated the lowest of all expressions in surprise and happiness, they were rated second only to anger faces in anger, a finding that is consistent with other evidence concerning similar reactions to neutral and anger expressions (e.g., Vrana, 2004).

Smile ratings of generalization faces

Although we selected the generalization faces to include only neutral expressions, there was some variability in smile ratings, and we therefore deleted all faces that received a mean rating greater than 2 on the 5-point smile scale, where 1 = *no smile*. This left a total of 312 faces. A 3 (face race) $\times$ 2 (face sex) analysis of variance (ANOVA) on the smile ratings for these faces revealed no significant effect of face sex ($M = 1.39$, $SD = 0.32$ and $M = 1.43$, $SD = 0.32$, for men and women, respectively), $F(1, 306) = 2.07$, $p = .15$, and no significant Face Race $\times$ Face Sex interaction, $F(2, 306) = 1.02$, $p = .36$. Although there was an effect of face race that approached significance, $F(2, 272) = 3.01$, $p < .06$, the mean ratings for all three groups were close to the low anchor on the scale. Post hoc Sidak tests³ revealed that Korean faces ($M = 1.42$, $SD = 0.30$) did not differ from either White faces ($M = 1.35$, $SD = 0.30$, $p = .28$) or Black faces ($M = 1.45$, $SD = 0.36$, $p = .88$), whereas White faces had marginally

³The Šidák correction is more precise than the well-known Bonferroni correction that closely approximates the Šidák equation. Although the latter is difficult to compute by hand, it is a post hoc test option in SPSS.

lower smile ratings than Black faces ($p=.08$). The final set of 312 faces (165 male) consisted of 109 White faces (57 male), 113 Korean faces (56 male), and 90 Black faces (52 male).

Emotion network activation by neutral faces as a function of sex and race

Mean activations of network units trained to recognize angry, happy, and surprise faces by the neutral generalization faces of each sex and race are shown in Table 1.

Surprise unit activation—As predicted, neutral expression female faces activated the network unit trained to recognize surprise faces significantly more than did male faces, $F(1, 306) = 12.38, p < .0001$. Face race also had a significant effect on activation of the surprise unit, $F(2, 306) = 122.26, p < .0001$. Post hoc Sidak tests revealed that Black faces activated the surprise unit more than did Korean faces ($p < .0001$), which in turn activated the surprise unit more than did White faces ($p < .0001$). The Face Race $\times$ Face Sex interaction was not significant, $F(2, 306) = 0.40, p = .67$.

Angry unit activation—As predicted, neutral expression male faces activated the network unit trained to recognize angry faces significantly more than did female faces, $F(1, 306) = 28.60, p < .0001$. A significant Face Race $\times$ Face Sex interaction, $F(2, 306) = 13.25, p < .0001$, further revealed that the tendency for male faces to activate the angry unit more than female faces was not significant for Korean faces ($p = .68$), whereas it was for both White and Black faces ($p < .0001$ and $p < .01$, respectively).

Face race also had a significant effect on activation of the angry unit, $F(2, 306) = 36.05, p < .0001$. Post hoc Sidak tests revealed that neutral expression White faces activated the angry unit more than did Korean faces ($p < .0001$) or Black faces ($p < .0001$), whereas the latter two did not differ from each other ($p = .72$). The significant Face Race $\times$ Face Sex interaction further revealed that the tendency for White faces to activate the angry unit more than Black or Korean faces was significant for male faces (both $ps < .0001$) and for White versus Black female faces ($p = .01$) but not for White versus Korean female faces ($p = .18$).

Happy unit activation—Contrary to prediction, neutral expression male faces activated the network unit trained to recognize happy faces significantly more than did female faces, $F(1, 306) = 10.25, p < .01$. Face race also had a significant effect on activation of the happy unit, $F(2, 306) = 67.25, p < .0001$. Post hoc Sidak tests revealed that neutral expression Black faces activated the happy unit significantly more than did White or Korean faces (both $ps < .0001$), which did not differ from each other ($p = .23$). The Face Race $\times$ Face Sex effect was not significant, $F(2, 306) = 0.15, p = .86$.

Relative activation of angry and happy units—Although male faces elicited greater activation of the output units trained to recognize either angry or happy faces, the sex difference was stronger for angry faces, consistent with prediction. We therefore examined sex differences in the relative activation of the two output units by subtracting happy unit activation from angry unit activation for each face (see Table 1). The ANOVA for this measure indicated that the relative activation of the angry versus happy unit was greater for male faces, $F(1, 306) = 7.09, p < .01$. A significant Face Race $\times$ Face Sex interaction effect, $F(2, 306) = 9.01, p < .0001$, further revealed that the face sex effect varied across face race. It was significant for White faces ($p < .0001$) but not for Korean faces ($p = .19$) or Black faces ($p = .21$).

There was also a main effect of face race, $F(2, 306) = 83.30, p < .0001$, and a Sidak post hoc test revealed that the relative activation of the angry versus happy unit was greater for White

faces than for Korean faces, which elicited greater relative activation of the angry unit than Black faces (all $ps < .0001$). Although the Sex $\times$ Race interaction revealed that the effects of face race varied as a function of face sex, the overall race differences were significant for both male and female faces (all $ps < .03$).

Discussion

Whereas both cultural stereotypes and structural commonalities could provide viable explanations for previous studies showing that sex and race moderated the recognition of angry, surprise, and happy emotion expressions and vice versa (Becker et al., 2007; Hess et al., 1997; Hugenberg & Bodenhausen, 2003, 2004; Le Gal & Bruce, 2002), Study 1 demonstrates similarities between emotion expressions and faces from various demographic categories that are due purely to structure. A connectionist model “knows” nothing but similarities in facial metrics and is uninfluenced by cultural stereotypes. The modeling revealed that the facial metrics of neutral expression female faces activated an output unit trained to recognize surprise expressions more than did those of male faces. This result supports a structural explanation for the moderation of surprise recognition by face sex (Le Gal & Bruce, 2002). It also is consistent with previous evidence that babyish faces resemble surprise expressions more than mature faces do (Zebrowitz, Kikuchi, & Fellous, 2007), because adult female faces retain more neotenous characteristics than do male faces. The modeling also revealed that the facial metrics of neutral expression male faces activated an output unit trained to recognize angry expressions more than did those of female faces, supporting a structural explanation for the moderation of anger recognition by face sex (Becker et al., 2007; Hess et al., 1997). Again the greater neoteny of female faces links this result with previous evidence that baby faces resemble angry expressions less than mature faces do (Marsh et al., 2005; Sacco & Hugenberg, 2009; Zebrowitz, Kikuchi, & Fellous, 2007). Finally, the modeling revealed a tendency for male faces to resemble happy expressions more than female faces do, which suggests that previous moderating effects of face sex on the response to happy expressions and vice versa reflect similarities in the cultural meaning of female faces and happy faces rather than similarities in their *structure*. Although it remains possible that there are facial qualities not included in the modeling that render female faces more structurally similar to happy ones, the metrics that we included were sufficient to communicate the meaning of happy expressions, as shown in Study 2, where resemblance to happy expressions had predictable effects on trait impressions. It is noteworthy that previous studies showing moderating effects of face sex required participants to simultaneously consider male and female angry and happy faces, whereas the connectionist models independently compared the extent to which male and female faces resembled angry versus neutral faces and the extent to which they resembled happy versus neutral faces. When we performed an analysis on the relative resemblance of men and women to angry versus happy faces, the results paralleled the previous research findings. Specifically, compared with female faces, male faces showed relatively greater resemblance to angry than to happy faces.

It is interesting that the greater resemblance of men's faces to anger expressions or to anger versus happy expressions did not hold true for Korean faces, whereas men's lesser resemblance to surprise expressions held true for faces of all races. These results are consistent with evidence that the sexual dimorphism that yields greater male resemblance to angry expressions is less pronounced in Asian than in other race faces (Brown, 1999; Kunjur, Sabesan, & Ilankovan, 2006). In contrast, the somewhat overlapping neotenous features that yield greater female resemblance to surprise expressions appear to be racially universal.

Although our assessment of whether faces of different races differ in their resemblance to emotion expressions was exploratory, the results are provocative. The facial metrics of neutral expression Black faces activated output units trained to respond to happy and surprise expressions significantly more than did those of Korean or White faces. One might attribute the lower activation by White faces to the fact that the network had been trained to differentiate neutral from emotional expressions in White faces. However, despite being trained on White faces, the anger output unit was activated more by White than by Korean or Black neutral expression generalization faces. It is noteworthy that the connectionist models were provided only with facial metric information. It would be interesting to determine whether skin tone information modulates racial differences in resemblance to emotions that are revealed in facial metrics. For example, the reddening of White faces in anger may cause darker skin to serve as a cue to anger in models trained on White faces, whereas the blanching of faces in fear, and perhaps surprise, may cause lighter skin to serve as a cue to fear or surprise. It also remains to be determined whether the observed racial differences in resemblance to emotion expressions hold true when models are trained on Black faces or Korean faces. Although research has documented the universality of emotion expression morphology across people from diverse groups, there is also some evidence for cultural differences (Elfenbein & Ambady, 2003), and it is possible that some differences would emerge with training on other race faces.

The fact that Black faces show more structural similarity to happy and surprise expressions and less similarity to anger expressions than White faces do suggest that recognition of faces as angry versus happy or surprised would be faster when White faces have angry expressions and Black faces have happy or surprise expressions than vice versa. In the case of happy and angry faces, this prediction runs counter to the prediction that perceivers will be faster to recognize a face when it has a stereotypically associated expression, a prediction supported by the finding that highly prejudiced White perceivers were more attuned to angry than to happy faces in Black faces, whereas this was not true for those low in prejudice (Hugenberg & Bodenhausen, 2003, 2004). However, as noted above, these studies did not use faces with naturalistic variations in race-related features, and it is possible that the results would be reversed if such faces were used. Alternatively, the effects of racial differences in objective facial resemblance to emotions and those of stereotypes about racial differences in behavioral resemblance to emotions might cancel each other out.

The lesser resemblance of Korean than White faces to angry expressions and their greater resemblance to surprise expressions is consistent with the higher babyfacedness of Korean faces (Zebrowitz, Bronstad, & Lee, 2007), because more babyish faces resemble anger less and surprise more than do mature faces (Marsh et al., 2005; Sacco & Hugenberg, 2009; Zebrowitz, Kikuchi, & Fellous, 2007). It may also reflect possible race differences in the relative vertical position of features, because it has been shown that shortening the vertical distance between the eyes and mouth increases perceptions of anger (Neth & Martinez, 2009). In addition, one might speculate that the fuller lips of neutral expression Black faces contribute to their greater resemblance to broadly smiling happy expressions and their lesser resemblance to more tight-lipped angry expressions. Future research would be useful to pinpoint the particular features that contribute to differences in the emotion resemblance of White, Black, and Korean faces. However, regardless of the explanation for the race differences, the fact that neutral expression White faces resemble angry faces more than do faces of other races, whereas neutral expression Black faces resemble happy or surprised expressions more, may have interesting implications for race-related trait impressions, a possibility that we investigate in Study 2.

Study 2: The Contribution of Resemblance to Emotion Expressions to Trait Impressions and Race Stereotypes

According to the emotion overgeneralization hypothesis (Zebrowitz, 1996, 1997), the adaptive value of responding appropriately to people in certain emotional states has produced a strong preparedness to respond to facial qualities that communicate emotion, with those responses overgeneralized to individuals whose facial structure merely resembles a particular emotion. Consistent with this hypothesis, Montepare and Dobish (2003) found that some neutral expression faces create perceptions of an angry demeanor and elicit impressions of low warmth and high dominance, whereas others create perceptions of a happy demeanor and elicit impressions of high warmth and high dominance. Moreover, these effects were shown with facial attractiveness and babyfacedness controlled, demonstrating that they held true over and above any effects of variations in facial maturity or attractiveness on resemblance to emotion. This control is important because, as noted above, babyish faces resemble anger less than do mature faces, and they also resemble surprise and fear expressions more (Marsh et al., 2005; Sacco & Hugenberg, 2009; Zebrowitz, Kikuchi, & Fellous, 2007). In addition, because smiling faces are judged more attractive (Reis et al., 1990), attractive faces might resemble happy expressions more.

Whereas Montepare and Dobish (2003) relied on perceivers' subjective judgments of resemblance to emotion expressions, which could be influenced by other aspects of the face that also influenced trait impressions, Said, Sebe, and Todorov (2009) assessed the objective resemblance of faces to various emotions with a Bayesian network classifier and found similar results: Neutral expression faces that were perceived as having warmer, less threatening traits showed a greater resemblance to happy or surprised faces, whereas neutral expression faces that were perceived as having more threatening traits showed a greater objective resemblance to angry faces. Similarly, Oosterhof and Todorov (2008) found that computer-generated neutral expression faces morphed to exaggerate features associated with high trustworthiness look happy, those morphed to exaggerate features associated with low trustworthiness look angry, and those morphed to exaggerate features associated with submissiveness look fearful. However, unlike the Montepare and Dobish (2003) study, these two studies did not control facial attractiveness and babyfacedness. Thus, the possible confounding of resemblance to surprise, anger, or happiness with variations in babyfacedness or attractiveness might contribute to their effects. In Study 2, we combined the strengths of the foregoing studies to assess whether an *objective* resemblance to emotion expressions predicts trait impressions over and above any prediction by attractiveness and babyfacedness. Specifically, we used connectionist modeling to provide an objective assessment of the extent to which neutral expression faces resembled surprise, happy, and angry expressions, and we statistically controlled ratings of attractiveness and babyfacedness when predicting trait impressions from emotion resemblance.

In addition to examining whether resemblance to emotion expressions contributes to trait impressions, we also examined in Study 2 whether the effects are moderated by face race and whether race differences in emotion resemblance contribute to race stereotypes. We could not address the contribution of emotion resemblance to sex stereotypes, because the experimental design was not sensitive to sex stereotypes. Judges rated either male or female faces, not both, thereby contrasting each face with others of the same sex. Consistent with a shifting standards effect (Biernat & Manis, 1994), this design did not yield significantly different trait impressions of male and female faces.

Our predictions concerning the contribution of emotion resemblance to race stereotypes derive from the race differences in resemblance to emotion expressions documented in Study 1, coupled with previous findings of White judges' racial stereotypes of the faces used

in Study 2. Specifically, White judges rated the Black faces as more dangerous, less competent, less likeable, and more expressive than the White faces, whereas they judged the Korean faces as less dangerous, more competent, less likeable, and less expressive than the White faces (Zebrowitz, Bronstad, & Lee, 2007). Because our judges were all White, we investigated the contribution to these race stereotypes of resemblance to emotion expressions using ingroup White faces as the standard of comparison.

We made the following predictions:

1. People whose neutral expression faces show greater structural resemblance to angry faces will elicit impressions associated with actual anger expressions: more dangerous, less likeable, and more expressive.
2. People whose neutral expression faces show greater structural resemblance to happy expressions will elicit impressions associated with actual happy expressions: less dangerous, more likeable, and more expressive.
3. People whose neutral expression faces show greater structural resemblance to surprise faces will elicit impressions associated with actual surprise expressions: less dangerous, less competent, and more expressive.
4. White judges' cultural stereotypes of Black targets as more dangerous and less likeable than White targets are partially suppressed by Black faces' weaker resemblance to angry expressions and their greater resemblance to happy expressions, as will be evidenced by the emergence of more negative impressions when resemblance to angry or happy expressions is statistically controlled.
5. White judges' cultural stereotypes of Korean targets as less dangerous than White targets are partially mediated by Korean faces' weaker resemblance to angry expressions, as will be evidenced by the weakening of this positive impression when resemblance to angry expressions is statistically controlled. However, impressions of Koreans as less likeable than White targets is suppressed by their weaker resemblance to angry expressions, as will be evidenced by the emergence of more negative impressions when resemblance to anger is controlled.

Method

Faces and network training

The facial metrics, connectionist model, network training, and network generalization were identical to Study 1.

Impressions of generalization faces

These measures are discussed in the sections below.

Judges—White judges rated random orders of either the male or female faces that were used as the generalization faces in Study 1. Ratings of 180 faces were taken from Zebrowitz, Bronstad, and Lee (2007), in which 20 White judges (10 men) rated 30 Korean and 30 White male faces; 20 White judges (10 men) rated 30 Korean and 30 White female faces; 20 White judges (10 men) rated 30 Black and 30 White male faces; 20 White judges (10 men) rated 30 Black and 30 White female faces. The White faces paired with the Korean or Black faces for the foregoing sets of judges were identical. For the present study, the two ratings were combined by computing the mean rating for each White face on each scale across all judges. Ratings of an additional 180 faces (60 of each race) were provided by 17 judges (7 men) rating 90 male faces (30 of each race); and 14 judges (7 men) rating 90 female faces (30 of each race).

Rating scales—Judges rated the faces on the following 7-point scales: *not at all likeable*–*very likeable*; *unintelligent*–*intelligent*; *not very hardworking*–*very hardworking*; *not at all hostile*–*very hostile*; *not at all trustworthy*–*very trustworthy*; *reserved/quiet*–*expressive/loud*; *not at all familiar*–*very familiar*; *babyfaced*–*maturefaced*; *unattractive*–*attractive*, with the latter two ratings used as control variables.⁴

Procedure—Judges viewed images and input responses on Pentium 4 personal computers with Windows XP and 19-in. (48-cm) CRT displays. MediaLab 2004.2.1 (Jarvis, 2004) was used to display images and collect ratings. Faces were displayed for 5 s, or until a rating was made. A random order of all faces (races mixed) was rated on one scale before judges proceeded to the next scale in order to minimize carryover effects in ratings from one scale to another. A different random order of faces was presented for each rating scale. All judges rated familiarity first, followed by likeability, with attractiveness and babyfacedness rated in counterbalanced order after the traits. The order of trait ratings was intelligent, expressive, hostile, trustworthy, and hardworking for half the judges and the reverse order for the other half.

Results

Reliability of trait ratings

Acceptable reliabilities for the judges rating either Black and White faces or Korean and White faces are reported in Zebrowitz, Bronstad, and Lee (2007). For judges who rated the additional faces of all races used in the present study, Cronbach alphas for male and female faces, respectively, were as follows: likeable α s = .82 and .84; intelligent α s = .86 and .79; hardworking α s = .88 and .69;⁵ hostile α s = .85 and .85; trustworthy α s = .85 and .81; expressive/loud α s = .83 and .87; attractive α s = .78 and .89; babyfaced α s = .85 and .80.

Trait composites

On the basis of the results of Zebrowitz, Bronstad, and Lee (2007), we created two a priori trait composite variables. Ratings of “hardworking” and “intelligent” were summed to create a *competence* composite, and we summed ratings of “hostile” and “trustworthy” (reverse scored) to create a *danger* composite. A factor analysis performed on ratings of the additional faces used in the present study confirmed these a priori composites, with ratings of hardworking and intelligent loading highest on the first factor, ratings of trustworthy and hostile loading highest on the second factor, and ratings of expressive/loud loading on a third factor (see Table 2). Likeability was not included in the factor analysis because it is conceptually distinct from personality traits and because Zebrowitz, Bronstad, and Lee (2007) found that low likeability co-occurred with positive trait impressions for Korean faces. Alpha coefficients for the competence and danger composites were .88 and .91, respectively. These two factors are congruent with decades of research showing that these are the two primary dimensions along which people evaluate others, although various factor labels have been used, such as *warmth* or *social goodness* as the positive pole for what we call *danger* (Cuddy & Fiske, 2002; Rosenberg, Nelson, & Vivekananthan, 1968; Wiggins, 1996).

⁴Ratings of familiar–unfamiliar were used as an additional control variable in a supplementary analysis summarized in Footnote 7.

⁵Six judges were dropped from Korean female hardworking ratings to achieve an acceptable reliability coefficient. The reliability for hardworking ratings of female faces was the average of those calculated for White (all judges), Black (all judges), and Korean (reduced judges) faces separately.

Overview of regression analyses

Given the high interjudge reliabilities, we computed mean ratings for each face across judges and used face as the unit of analysis. As in Study 1, we deleted faces with smile ratings > 2. We performed separate regression analyses to predict each of the impressions (competence, danger, likeability, and expressive/loud) from resemblance to angry expressions, happy expressions, and surprise expressions.

Our analyses examined three effects: (a) whether there is emotion expression overgeneralization on impressions independent of face race and other appearance qualities; (b) whether emotion expression overgeneralization contributes to (i.e., mediates or suppresses) racially stereotyped impressions; and (c) whether emotion overgeneralization differs (i.e., is moderated) by face race. In Step 1, we predicted impressions from face race as a set of dummy variables that treat White faces as a reference group, with face sex, attractiveness, and babyfacedness as control variables. The betas for face race at this step reveal our White judges' stereotypes of Korean faces and Black faces, which may reflect both the *cultural meaning* of the demographic categories and their *resemblance to emotion expressions*. In Step 2, we added the emotion unit trained to distinguish anger, happy, or surprise from neutral expressions. The beta for emotion unit activation at this step reveals whether there is an emotion overgeneralization effect on impressions independent of face race and other appearance qualities. We also examined the change in the beta for face race from Step 1 to Step 2, and we tested whether emotion overgeneralization mediates or suppresses racial stereotypes following the procedure outlined by Baron and Kenny (1986). For example, if the tendency to perceive Korean faces as less dangerous than White faces at Step 1 is significantly reduced when anger unit activation is controlled at Step 2, then that would indicate that Korean faces' lower resemblance to angry faces mediates the tendency to perceive them as less dangerous than White faces. Similarly, if the tendency to perceive Black faces as more dangerous than White faces at Step 1 is significantly increased when happy unit activation is controlled at Step 2, then that would indicate that Black faces' greater resemblance to happy faces suppresses the tendency to perceive them as more dangerous than White faces. Finally, in Step 3, we entered the interaction term of Emotion Unit Activation $\times$ Face Race. A significant beta for this interaction term would indicate that the effect of resemblance to emotion expressions on trait impressions was moderated by face race. To elucidate significant interaction effects, we performed regression analyses predicting trait impressions within face race, controlling the main effects for face race and emotion unit activation as well as face sex, attractiveness, and babyfacedness.

Effects of face race on trait impressions

Step 1 in the regression analyses shows the effects of face race on trait impressions and likeability, controlling for face sex, attractiveness, and babyfacedness.⁶ Compared with White faces, Korean faces were rated as more competent, less dangerous, less expressive, and less likeable, and Black faces were rated as less competent, more dangerous, and less likeable (see Table 3). These effects replicate the results of Zebrowitz, Bronstad, and Lee (2007) with a larger set of faces, and we do not discuss them further.

Effects of resemblance to angry expressions

The effects are discussed below.

⁶The effects of the control variables are not shown in Tables 4, 5, and 6. More attractive faces were judged more competent, less dangerous, more likeable, and more expressive/loud; more babyfaced were judged more competent (an atypical effect that can be attributed to the fact that the Korean faces were more babyfaced and also judged more competent), less dangerous, more likeable, less expressive/loud; female faces were judged less likeable.

Overall effects and moderation by face race—As predicted, neutral expression faces that elicited greater activation of the angry unit entered in Step 2 were judged as more dangerous and less likeable (see Table 4). They were also judged as less competent. All of these effects were qualified by Black versus White Face Race $\times$ Anger Unit Activation interaction effects in Step 3. Regressions within face race revealed that the negative effect of anger unit activation on ratings of competence was shown for White faces ($\beta = -.23$, $p = .03$), but not for Black faces ($\beta = -.03$, $p = .73$), which were judged low in competence regardless of degree of resemblance to angry faces (see Figure 3A). The positive effect of anger unit activation on ratings of danger and the negative effect on ratings of likeability also were shown for White faces (respective β s = .18 and $-.13$, p s = .06 and .07) but not for Black faces (respective β s = .06 and $-.02$, p s > .55), which again were judged high in danger and low in likeability regardless of resemblance to angry faces (see Figures 3B and 3C). The marginally significant interaction effect on impressions of expressivity reflected a positive influence for White faces and a negative influence for Black faces, neither of which were significant (respective β s = .09 and $-.06$, p s > .34) (see Figure 3D). Finally, White judges showed sensitivity to variations in resemblance to anger among Korean faces in all their impressions, with the simple effects of anger unit activation significant for impressions of Koreans' competence, danger, likeability, and expressivity (β s = $-.21$, .21, $-.14$, and .24; all p s < .01, except likeability, $p = .03$) (see Figures 3A–3D).

Mediation of race-related impressions by resemblance to angry expressions

The tendency for Korean faces to be rated as more competent and less dangerous than White faces was weakened when anger node activation was entered in Step 2, and anger node activation qualified as a potential mediator of these effects. Sobel tests revealed that the tendency for neutral expression Korean faces to activate the anger node less than neutral expression White faces mediated the tendency to perceive Korean faces as more competent and less dangerous than White faces (Sobel test z s = 3.50 and 3.69, respectively; p s < .001). In contrast, the tendency for neutral expression Black faces to activate the anger node less than neutral expression White faces partially suppressed the tendency to perceive Black faces as less competent and more dangerous than White faces (Sobel test z s = 3.37 and 3.54, p s < .001). Finally, the perception of both Korean and Black faces as less likeable than White faces was partially suppressed by their lesser activation of the anger unit (Sobel test z s = 3.08 and 2.99, respectively; both p s < .01).

Effects of resemblance to happy expressions

The effects are discussed below.

Overall effects and moderation by face race—As predicted, the neutral expression faces that elicited greater activation of the happy unit entered in Step 2 were perceived as less dangerous, more likeable, and marginally more expressive (see Table 5). However, the effect for likeability was qualified by a significant Korean versus White Face Race $\times$ Happy Unit activation in Step 3. Regressions within face race revealed that Korean faces with a greater resemblance to happy expressions tended to be judged more likeable ($\beta = .16$, $p = .02$), whereas the effect was not significant for own-race White faces ($\beta = .02$, $p = .71$). Although the Black versus White Face Race $\times$ Happy Unit activation was not significant, regressions within face race revealed that the simple effect for Black faces approached significance ($\beta = .14$, $p = .06$) (see Figure 3E).

Mediation of race-related impressions by resemblance to happy expressions

—The tendency for Black faces to be rated as more dangerous and less likeable than White faces in Step 1 were both strengthened when happy node activation was entered in Step 2, and happy node activation qualified as a potential mediator of these effects. Sobel tests

revealed that the tendency for neutral expression Black faces to activate the happy node more than the neutral White faces was a partial suppressor of the tendency to perceive Black faces as more dangerous and less likeable than White faces (Sobel test z s = 2.54 and 2.40, respectively; p s < .02).

Effects of resemblance to surprise expressions

The extent to which neutral expression faces activated the surprise node in Step 2 did not influence impressions, and there were no significant interactions with face race (see Table 6). Because surprise expressions are structurally similar to a babyish appearance (Zebrowitz, Kikuchi, and Fellous, 2007), we ran a supplementary set of regression analyses to determine whether resemblance to surprise would influence additional trait impressions when babyfacedness was not included as a control variable, thereby retaining any overlapping variance of babyfacedness and resemblance to surprise. Although there were no main effects of surprise activation on impressions and no interactions with face race, regressions within face race revealed that Korean faces that activated the surprise node more were perceived as less dangerous ($\beta = -.19$, $p = .03$) and marginally more likeable ($\beta = .13$, $p = .06$).⁷

Discussion

Consistent with the emotion overgeneralization hypothesis, Study 2 revealed that trait impressions of neutral expression faces are influenced by their structural resemblance to emotion expressions. Neutral expression faces that showed greater resemblance to anger expressions were perceived as higher in danger (more hostile and less trustworthy), lower in perceived likeability, and lower in competence (less hardworking and intelligent). Whereas the inference of dangerous behavioral predispositions from facial resemblance to anger were predicted as a logical extension of responses to actual angry faces, the inference of low competence suggests that resemblance to anger produces a more general negative halo effect even though attractiveness as well as babyfacedness were statistically controlled. The fact that resemblance to anger had predictable effects on impressions indicates that the model adequately captured resemblance to anger in Study 1 even though, as is also true for human judges, accuracy in differentiating anger from neutral expressions was lower than accuracy in differentiating happy or surprise from neutral expressions. In contrast to faces that resembled anger, neutral expression faces that showed greater resemblance to happy expressions elicited impressions of less danger, more likeability, and more expressivity, logical extensions of reactions to happy people, whereas resemblance to happy faces did not show a more generalized positive halo effect on impressions of competence. Although impressions of neutral expression faces were not influenced by their resemblance to surprise, exploratory analyses revealed that when babyfacedness was not controlled, Korean faces with more resemblance to surprise were perceived as less dangerous and more likeable.

Although this is the first evidence that resemblance to anger decreases a person's perceived competence, the effects of resemblance to angry and happy expressions on impressions of

⁷Because Zebrowitz, Bronstad, and Lee (2007) had shown that the lesser familiarity of other-race faces mediated negative stereotypes and suppressed positive ones, and because familiar faces are judged happier and less angry than novel ones (Claypool, Hugenberg, Housley, & Mackie, 2007), familiarity ratings were used as an additional control variable in supplementary analyses to ensure that the mediation and suppression effects produced by resemblance to emotion expressions held true even with familiarity controlled. Consistent with previously documented effects, controlling the lesser familiarity of other-race faces produced some differences in Step 1 of the regressions, such that Black faces were no longer rated as less competent than White faces, and Black and Korean faces were no longer rated as less likeable than White faces. However, the pattern of mediation/suppression effects for resemblance to emotion expressions remained the same. For example, adding angry or happy unit activation in Step 2 yielded a significant difference in ratings of the likeability of Black and White faces that was not seen in Step 1 with familiarity controlled, demonstrating an added suppressor effect of Black faces lesser resemblance to angry expressions and greater resemblance to happy expressions. Likewise, Korean faces were rated as less likeable at Step 2 when their lesser resemblance to anger was controlled, again demonstrating that their lesser resemblance to angry expressions partially suppressed impressions of low likeability.

danger and likeability are consistent with previous evidence that neutral expression faces judged to resemble anger elicit impressions of low warmth, whereas those judged to resemble happiness elicit impressions of high warmth (Montepare & Dobish, 2003); those with an objectively greater resemblance to angry or happy faces elicit impressions of greater or lesser threat, respectively (Said et al., 2009), and those morphed to exaggerate features associated with trustworthiness or untrustworthiness look happy or angry, respectively (Oosterhof & Todorov, 2008). Our findings strengthen the conclusions of past research by showing that the effects hold true given both an objective measure of resemblance to happy and angry expressions as well as statistical control of attractiveness and babyfacedness. This demonstrates that emotion overgeneralization effects are independent of an attractiveness halo effect (anomalous face overgeneralization) and babyface overgeneralization (Zebrowitz, Fellous, Mignault, & Andreoletti, 2003). The effects were also independent of familiar face over-generalization effects (Zebrowitz, Bronstad, & Lee, 2007; see Footnote 7).

Whereas our findings for resemblance to anger and happy expressions replicated effects reported by Said et al. (2009) even when we controlled attractiveness and babyfacedness, we did not replicate their finding that neutral expression faces perceived as having warmer, less threatening traits showed a greater resemblance to surprised faces. However, when we removed babyfacedness as a control variable, we then found that resemblance to surprise predicted impressions of lower danger and likeability, albeit only in Korean faces. These results suggest that the surprise expression overgeneralization effect for impressions of threat may be redundant with a babyface overgeneralization effect. Surprise faces and baby's faces share the feature of relatively large eyes, both of which elicit impressions of warmer, less threatening traits (Zebrowitz, Kikuchi, & Fellous, 2007). The fact that the effects of surprise were restricted to Korean faces when babyfacedness was not controlled may be due to the fact that Asian faces have relatively smaller eyes than the other two groups, providing more opportunity for within-group variations in eye size that increase resemblance to surprise to influence impressions.

Additional moderating influences of face race suggest that the effects of resemblance to anger and happy expressions depend on the baseline level of impressions. Thus, greater resemblance to anger decreased the White judges' impressions of the competence and likeability of White and Korean faces and increased impressions of their danger, whereas it did not decrease the already low judged competence and likeability of Black faces or increase their already high judged danger. Similarly, greater resemblance to happy expressions influenced the likeability of Korean and Black faces, but not of White faces, which were more well liked to begin with. It seems reasonable to argue that these moderation effects reflect White judges' cultural stereotypes rather than racial variations in facial structure. For example, even though both Black and Korean faces resembled angry expressions less than White faces did, within-group variations in Black and Korean resemblance to angry faces had divergent effects on impressions. Similarly, even though Black but not Korean faces resembled happy expressions more than White faces did, within-group variations in Black and Korean resemblance to happy expressions had similar effects on impressions. Insofar as these moderation effects do reflect the cultural meaning of the faces to our White judges, the pattern of results should be different for judges of other races. For example, because Black judges rate Black faces as more likeable than White ones (Zebrowitz, Bronstad, & Lee, 2007), resemblance to anger may decrease Black judges' impressions of the likeability of Black faces more than White faces, and resemblance to happy expressions may increase Black judges' impressions of the likeability of White faces more than Black faces.

In addition to the foregoing direct effects of resemblance to an emotion expression on first impressions of faces and their moderation by face race, Study 2 showed that such resemblance both mediated and suppressed White judges' race stereotypes. White judges' impressions of Black faces as more dangerous and less likeable than White faces were partially suppressed by Black faces' greater resemblance to happy faces, and these impressions as well as impressions of lower competence were also partially suppressed by Black faces' lesser resemblance to angry faces. Similarly, White judges' impressions of Korean faces as less dangerous and more competent than White faces, effects consistent with the "model minority" stereotype (Kawai, 2005; Lin, Kwan, Cheung, & Fiske, 2005), were mediated by their lesser resemblance to angry faces, whereas their impressions of Korean faces as less likeable than White faces were partially suppressed by their lesser resemblance to angry faces.

In summary, Study 2 supports four general conclusions regarding emotion overgeneralization effects. First, an objective resemblance of neutral expression faces to emotion expressions influences trait impressions over and above any influence of facial attractiveness or babyfacedness. Second, emotion resemblance can influence not only trait impressions conceptually relevant to the emotion but also impressions that are merely similarly valenced to the emotion. Third, angry emotion overgeneralization effects are more likely when the target belongs to a social category viewed *positively* on the attribute being rated, whereas happy emotion overgeneralization is more likely to influence impressions when the target belongs to a category viewed negatively on the attribute being rated. Finally, emotion overgeneralization not only contributes to first impressions of neutral expression faces, but it also mediates and suppresses stereotypes of social categories that differ in their resemblance to emotion expressions.

General Discussion

The present research is the first to demonstrate *objective* similarities between the appearance of certain emotion expressions and faces that vary in sex or race. Consistent with previous research on the speed and accuracy of recognizing surprise and angry expressions in faces that vary in sex, neutral expression male faces showed objectively greater similarity to angry expressions as compared with female faces, which showed greater similarity to surprise faces. These results reveal that one need not invoke similar expectations about angry emotions and men or surprise emotions and women to explain the moderating effects of face emotion on sex recognition and vice versa. Rather, the greater ease of recognizing anger in men than in women and of recognizing surprise in women than in men can be explained by intrinsic structural properties of the faces, although differences in the cultural meaning of female and male faces could certainly add to the effects. However, the greater speed and accuracy of recognizing happy expressions in women than in men seems to reflect similarities in the cultural meaning of the faces, because it cannot be explained by a greater physical resemblance between female and happy faces.

Faces of different races, like those of different sexes, were shown to vary in their objective resemblance to emotion expressions. Neutral expression White faces showed greater resemblance to angry expressions than did Black or Korean faces, whereas neutral expression Black faces showed greater resemblance to happy and surprise expressions than did White faces. Because the emotion faces on which the connectionist models were trained were all White, it remains to be determined whether these race differences in resemblance to emotion expressions hold true when models are trained on Black faces or on Korean faces. Race differences in resemblance to emotion expressions might also vary if skin tone information were provided to the models in addition to facial metrics. Finally, it should be acknowledged that there is considerable variability within what we have called White,

Black, and Asian faces. For example, there are structural differences between the East Asian Korean faces that we used in the present study and South Asian faces, which could have implications for resemblance to emotion expressions. Similarly, differences between our sample of African American faces and some East African faces as well as differences between samples of White faces representing different ethnicities could yield differences in resemblance to emotion expressions.

The present research is also the first to demonstrate that resemblance to emotion expressions influences trait impressions both when emotion resemblance is objectively assessed and when possible confounding qualities of attractiveness and babyfacedness are controlled. Some of these effects were moderated by face race in ways that suggest that the pattern of results may be different for judges of other races. Nevertheless, the emotion overgeneralization effects that were shown in impressions of other-race faces reveal a fine-grained response that belies the view that “they all look alike,” a finding that is consistent with other evidence that reactions to other-race individuals are not simply categorical (Blair, Judd, & Chapleau 2004; Blair et al., 2002; Eberhart, Davies, Purdie-Vaughns, & Johnson, 2006; Livingston & Brewer, 2002; Maddox, 2004; Zebrowitz, Montepare, & Lee, 1993).

Finally, the present research is the first to demonstrate that racial differences in emotion resemblance contribute to race stereotypes. The lesser resemblance to angry faces of Black or Korean than White neutral expression faces suppressed White judges’ negative stereotypes of African Americans and mediated their positive stereotypes of Asians, and the greater resemblance of Black faces to happy expressions had parallel effects. We did not investigate the contribution of sex differences in emotion resemblance to sex stereotypes, because the between-groups rating of male and female faces yielded no significant differences. However, it would be interesting to conduct a conceptual replication of the present research to see whether the greater resemblance of men to angry versus to happy faces mediates the tendency to stereotype men as more threatening than women and whether the greater resemblance of women to surprise faces mediates the tendency to stereotype women as less threatening than men.

In summary, the emotion overgeneralization effects that we have documented may contribute to biases in correctly identifying the emotions of people who vary in sex or race as well as to biased first impressions. They have implications for intergroup perception and interaction not only among ordinary citizens, who increasingly interact with people of diverse racial backgrounds in an ever shrinking world, but even in international relations. For example, the lesser resemblance of Korean than White faces to angry expressions could contribute to a failure by White diplomats to recognize expressions of anger by their Asian counterparts, to underestimate their threat, and to overestimate their competence. Similarly, if the present findings generalize to judges of other races, then the greater resemblance of White than Black or Korean faces to angry expressions could contribute to a tendency for Asian or African diplomats to misconstrue their White counterparts’ neutral expressions as anger, to overestimate their threat, and to underestimate their competence. Thus, smooth intergroup relations may be disrupted not only by invidious outgroup stereotypes but also by adaptive responses to emotion expressions that are over-generalized to groups whose faces subtly resemble particular emotions.

Acknowledgments

This work was supported by National Institutes of Health Grants MH066836 and K02MH72603 awarded to Leslie A. Zebrowitz. We thank Alain Mignault for his work on the connectionist modeling code; Irene Blair for sharing her images of Black male faces with us; and Dena Leeman, Aviva Androphy, Shayna Skelley, Naomi Skop, Rebecca Wiernerman, and Glenn Landauer for their help finding, scanning, and standardizing facial images and collecting data.

References

- Baron RM, Kenny DA. The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*. 1986; 57:1173–1182. [PubMed: 3806354]
- Becker DV, Kenrick DT, Neuberg SL, Blackwell KC, Smith DM. The confounded nature of angry men and happy women. *Journal of Personality and Social Psychology*. 2007; 92:179–190. [PubMed: 17279844]
- Biernat M, Manis M. Shifting standards and stereotype-based judgments. *Journal of Personality and Social Psychology*. 1994; 66:5–20. [PubMed: 8126651]
- Blair I, Judd CM, Chapleau KM. The influence of Afrocentric facial features in criminal sentencing. *Psychological Science*. 2004; 15:674–679. [PubMed: 15447638]
- Blair IV, Judd CM, Sadler MS, Jenkins C. The role of Afrocentric features in person perception: Judging by features and categories. *Journal of Personality and Social Psychology*. 2002; 83:5–25. [PubMed: 12088132]
- Brown T. Facial profiles: Population standards and variability. *Perspectives in Human Biology*. 1999; 4:135–144.
- Claypool HM, Hugenberg K, Housley MK, Mackie DM. Familiar eyes are smiling: On the role of familiarity in the perception of facial affect. *European Journal of Social Psychology*. 2007; 37:856–866.
- Cuddy, AJC.; Fiske, ST. Doddering but dear: Content and functioning in stereotyping of older persons. In: Nelson, TD., editor. *Ageism: Stereotyping and prejudice against older people*. Cambridge, MA: MIT Press; 2002. p. 3–26.
- Eberhardt JL, Davies PG, Purdie-Vaughns VJ, Johnson SL. Looking deathworthy: Perceived stereotypicality of Black defendants predicts capital-entencing outcomes. *Psychological Science*. 2006; 17:383–386. [PubMed: 16683924]
- Elfenbein HA, Ambady N. Universals and cultural differences in recognizing emotions of a different cultural group. *Current Directions in Psychological Science*. 2003; 12:159–164.
- Enlow, DH. *Facial growth*. 3. Philadelphia, PA: Harcourt Brace; 1990.
- Hess U, Blairy S, Kleck RE. The intensity of emotional facial expressions and decoding accuracy. *Journal of Nonverbal Behavior*. 1997; 21:241–257.
- Hess U, Blairy S, Kleck RE. The influence of facial emotion displays, gender, and ethnicity on judgments of dominance and affiliation. *Journal of Nonverbal Behavior*. 2000; 24:265–283.
- Hugenberg K, Bodenhausen GV. Facing prejudice: Implicit prejudice and the perception of facial threat. *Psychological Science*. 2003; 14:640–643. [PubMed: 14629699]
- Hugenberg K, Bodenhausen GV. Ambiguity in social categorization: The role of prejudice and facial affect in race categorization. *Psychological Science*. 2004; 15:342–345. [PubMed: 15102145]
- Jarvis, BG. *Medialab* (Version 2004, 3.24) [Computer software]. New York, NY: Empirisoft Corporation; 2004.
- Kawai Y. Stereotyping Asian Americans: The dialectic of the model minority and the yellow peril. *Howard Journal of Communications*. 2005; 16:109–130.
- Knutson B. Facial expressions of emotion influence interpersonal trait inferences. *Journal of Nonverbal Behavior*. 1996; 20:165–182.
- Kunjur J, Sabesan T, Ilankovan V. Anthropometric analysis of eyebrows and eyelids: An inter-racial study. *British Journal of Oral and Maxillofacial Surgery*. 2006; 44:89–93. [PubMed: 15936856]
- Le Gal PM, Bruce V. Evaluating the independence of sex and expression in judgments of faces. *Perception & Psychophysics*. 2002; 64:230–243. [PubMed: 12013378]
- Lewontin, R. Invited talk in lecture series on Biology, race, and genetics: Old problems and new findings. Brandeis University; 2001. Is race a sensible genetic concept?.
- Lin MH, Kwan VSY, Cheung A, Fiske ST. Stereotype content model explains prejudice for an envied outgroup: Scale of Anti-Asian American stereotypes. *Personality and Social Psychology Bulletin*. 2005; 31:34–47. [PubMed: 15574660]

- Livingston RW, Brewer MB. What are we *really* priming? Cue-based versus category-based processing of facial stimuli. *Journal of Personality and Social Psychology*. 2002; 82:5–18. [PubMed: 11811634]
- Maddox KB. Perspectives on racial phenotypicality bias. *Personality and Social Psychology Review*. 2004; 8:383–401. [PubMed: 15582860]
- Marsh AA, Adams RB Jr, Kleck RE. Why do fear and anger look the way they do? Form and social function in facial expressions. *Personality and Social Psychology Bulletin*. 2005; 31:73–86. [PubMed: 15574663]
- Martinez, AM.; Benavente, R. The AR face database. 1998 Jun. (CVC Report No.24). Retrieved from [www.http://cobweb.ecn.purdue.edu/~aleix/aleix_face_DB.html](http://cobweb.ecn.purdue.edu/~aleix/aleix_face_DB.html)
- Montepare JM, Dobish H. The contribution of emotion perceptions and their overgeneralizations to trait impressions. *Journal of Nonverbal Behavior*. 2003; 27:237–254.
- Neth D, Martinez AM. Emotion perception in emotionless face images suggests a norm-based representation. *Journal of Vision*. 2009; 9:1–11. [PubMed: 19271875]
- Oosterhof NN, Todorov A. The functional basis of face evaluation. *Proceedings of the National Academy of Sciences, USA*. 2008; 105:11087–11092.
- Reis HT, Wilson IM, Monestere C, Bernstein S, Clark K, Seidl E, et al. What is smiling is beautiful and good. *European Journal of Social Psychology*. 1990; 20:259–267.
- Rosenberg S, Nelson C, Vivekananthan PS. A multidimensional approach to the structure of personality impressions. *Journal of Personality and Social Psychology*. 1968; 9:283–294. [PubMed: 5670821]
- Sacco DF, Hugenberg K. The look of fear and anger: Facial maturity modulates recognition of fearful and angry expressions. *Emotion*. 2009; 9:39–49. [PubMed: 19186915]
- Said CP, Sebe N, Todorov A. Structural resemblance to emotional expressions predicts evaluation of emotionally neutral faces. *Emotion*. 2009; 9:260–264. [PubMed: 19348537]
- Secord, P. Facial features and inference processes in interpersonal perception. In: Tagiuri, R.; Petrullo, L., editors. *Person perception and interpersonal behavior*. Stanford, CA: Stanford University Press; 1958. p. 300–315.
- Vrana SR. Reactions to facial expressions: Effects of social context and speech anxiety on responses to neutral, anger, and joy expressions. *Biological Psychology*. 2004; 66:63–78. [PubMed: 15019171]
- Wiggins JS. An informal history of the interpersonal circumplex tradition. *Journal of Personality Assessment*. 1996; 66:217–233. [PubMed: 16367700]
- Zebrowitz, LA. Physical appearance as a basis for stereotyping. In: McRae, MHN.; Stangor, C., editors. *Foundation of stereotypes and stereotyping*. New York, NY: Guilford Press; 1996. p. 79–120.
- Zebrowitz, LA. *Reading faces: Window to the soul?*. Boulder, CO: Westview Press; 1997.
- Zebrowitz LA, Bronstad PM, Lee HK. The contribution of face familiarity to ingroup favoritism and stereotyping. *Social Cognition*. 2007; 25:306–338.
- Zebrowitz LA, Fellous JM, Mignault A, Andreoletti C. Trait impressions as overgeneralized responses to adaptively significant facial qualities: Evidence from connectionist modeling. *Personality and Social Psychology Review*. 2003; 7:194–215. [PubMed: 12788687]
- Zebrowitz LA, Kikuchi M, Fellous JM. Are effects of emotion expression on trait impressions mediated by babyfacedness? Evidence from connectionist modeling. *Personality and Social Psychology Bulletin*. 2007; 33:648–662. [PubMed: 17440203]
- Zebrowitz LA, Montepare JM, Lee HK. They don't all look alike: Differentiating same versus other race individuals. *Journal of Personality and Social Psychology*. 1993; 65:85–101. [PubMed: 8355144]
- Zuckerman M. Some dubious premises in research and theory on racial differences. *American Psychologist*. 1990; 45:1297–1303. [PubMed: 2285179]

Figure 1.
Location of points. When identical points are marked on the right and left sides, only those on the person's right side are indicated.

Figure 2. Location of facial metrics used as inputs to the connectionist models trained on facial metrics. All metrics were normed by E2, interpupillary distance.

Figure 3.

The effect of neutral expression faces' emotion unit activation on trait impressions as a function of face race. Trait impressions are controlled for face sex and mean ratings of attractiveness and babyfacedness. Centered values of emotion unit activation are shown at the mean and one standard deviation above and below the mean. A: The effect of neutral expression faces' Angry unit activation on impressions of competence. B: The effect of neutral expression faces' Angry unit activation on impressions of danger. C: The effect of neutral expression faces' Angry unit activation on impressions of likeability. D: The effect of neutral expression faces' Angry unit activation on impressions of expressive/loud. E: The effect of neutral expression faces' Happy unit activation on impressions of likeability.

Table 1

Emotion Unit Activation as a Function of Face Sex and Race

Face race	Face sex	N	Surprise		Angry		Happy		Angry-Happy	
			M	SD	M	SD	M	SD	M	SD
All	Male	165	55.03 ^a	22.34	45.66 ^a	24.67	39.11 ^a	18.32	6.55 ^a	31.64
	Female	147	60.30 ^b	22.94	33.44 ^b	21.84	32.67 ^b	16.68	0.77 ^b	26.89
White	Both	109	41.22 ^a	19.85	53.40 ^a	25.01	28.09 ^a	14.81	25.31 ^a	26.57
Korean		113	56.43 ^b	16.61	31.44 ^b	18.51	31.52 ^a	13.45	-0.08 ^b	20.44
Black		90	78.60 ^c	14.32	34.18 ^b	22.30	51.47 ^b	16.51	-17.30 ^c	25.34
White	Male	57	39.03 ^a	18.42	66.10 ^a	18.81	30.84 ^a	16.21	35.26 ^a	24.48
Korean		56	52.91 ^b	17.25	30.63 ^b	16.72	33.60 ^a	13.03	-2.97 ^b	19.15
Black		52	74.84 ^c	14.83	39.44 ^b	22.66	54.11 ^b	16.37	-14.67 ^c	26.20
White	Female	52	43.61 ^a	21.23	39.48 ^a	23.64	25.08 ^a	12.59	14.41 ^a	24.60
Korean		57	59.89 ^b	15.33	32.24 ^{a,b}	20.23	29.47 ^a	13.66	2.77 ^b	21.43
Black		38	83.75 ^c	11.95	26.98 ^b	19.91	47.87 ^b	16.22	-20.89 ^c	23.99
All faces		312	57.51	22.74	39.90	24.13	36.08	17.83	3.83	29.60

Note. Means within a group that have different superscripts differ at $p < .0001$ with Sidak correction except for the following: Korean and Black male faces differ at $p = .075$, and White and Black female faces differ at $p = .013$ for angry unit activation; male and female faces across face race differ at $p = .002$ for happy unit activation; male and female faces across face race differ at $p = .008$. Korean and Black male faces differ at $p = .029$, and White and Korean female faces differ at $p = .029$ for the difference score of angry minus happy unit activation.

Table 2

Factor Loadings of Trait Ratings

Trait	Factor		
	Competence	Danger	Expressive/loud
Hardworking	.87	.34	-.25
Intelligent	.78	.51	-.09
Hostile	-.35	-.90	.18
Trustworthy	.58	.75	-.09
Expressive/loud	-.15	-.13	.98

Note. Factors solutions were generated by varimax rotation with the number of factors specified as 3. Traits that were included in each factor are in bold.

Table 3

Mean Ratings of Neutral Expression Faces of Each Race Controlling for Attractiveness, Babyfacedness, and Face Sex

Face race	N	Competence		Danger		Likeable		Expressive/loud	
		M	SE	M	SE	M	SE	M	SE
White	109	4.14 ^a	.05	3.81 ^a	.06	3.86 ^a	.05	3.95 ^a	.05
Korean	113	4.75 ^b	.05	3.51 ^b	.06	3.72 ^b	.05	3.12 ^b	.05
Black	90	3.93 ^c	.05	4.17 ^c	.06	3.63 ^b	.05	4.05 ^a	.06

Note. Means were derived from an analysis of covariance, controlling attractiveness, babyfacedness, and face sex to correspond to the regression analyses. Attractiveness did not vary across face race, $F(2, 309) = 0.97$, $p = .38$, with mean ratings for White, Korean, and Black faces being 3.33 ($SD = 0.91$), 3.44 ($SD = 0.63$), and 3.49 ($SD = 0.93$). Babyfacedness did vary with race, $F(2, 309) = 26.32$, $p < .0001$, and Korean faces ($M = 4.28$, $SD = 0.73$) were rated more babyfaced than either Black ($M = 3.67$, $SD = 0.93$) or White faces ($M = 3.45$, $SD = 0.97$; $ps < .0001$), which did not differ. Means with different superscripts differ at $p < .001$ or better, except for the difference between competence ratings for White and Black faces at $p = .002$ and the difference between likeability ratings for White and Korean faces at $p = .046$.

Table 4

Predicting Impressions of Generalization Faces From Activation of the Angry Output Unit Controlling for Attractiveness, Babyfacedness, and Face Sex

Regression step	Predictor	Competence			Danger			Likeable			Expressive/loud		
		B	(SE)	β	B	(SE)	β	B	(SE)	β	B	(SE)	β
Step 1	Intercept	2.90	(.163)		6.04	(.193)		1.46	(.160)		3.31	(.184)	
	Korean (vs. White faces)	0.61	(.070)	.45***	-0.30	(.083)	-.19***	-0.14	(.069)	-.09*	-0.82	(.079)	-.50***
	Black (vs. White faces)	-0.21	(.069)	-.15**	0.37	(.082)	.22***	-0.23	(.068)	-.15**	0.11	(.078)	.06
Step 2	Intercept	2.92	(.159)		6.01	(.188)		1.48	(.158)		3.30	(.185)	
	Korean (vs. White faces)	0.48	(.075)	.36***	-0.15	(.089)	-.09	-0.24	(.075)	-.17**	-0.79	(.087)	-.48***
	Black (vs. White faces)	-0.32	(.073)	-.22***	0.50	(.086)	.30***	-0.32	(.072)	-.20***	0.14	(.084)	.08
	Angry unit activation	-0.01	(.001)	-.18***	0.01	(.002)	.20***	-0.004	(.001)	-.14**	0.001	(.001)	.04
Step 3	Intercept	2.93	(.159)		5.99	(.189)		1.50	(.158)		3.27	(.184)	
	Korean (vs. White faces)	0.45	(.077)	.33***	-0.11	(.091)	-.07	-0.27	(.076)	-.18***	-0.76	(.089)	-.46***
	Black (vs. White faces)	-0.34	(.075)	-.23***	0.50	(.089)	.30***	-0.32	(.074)	-.21***	0.11	(.087)	.07
	Angry unit activation	-0.01	(.002)	-.28***	0.01	(.002)	.26**	-0.01	(.002)	-.21**	0.002	(.002)	.07
	Angry $\times$ Korean	0.003	(.003)	.06	-0.001	(.004)	-.01	0.001	(.003)	.02	0.003	(.004)	.05
	Angry $\times$ Black	0.01	(.003)	.11*	-0.01	(.003)	-.09	0.01	(.003)	.08 [†]	-.01	(.003)	-.09 [†]

Note. Face sex was coded 0 for male faces and 1 for female faces. Korean was coded 1 for Korean faces and 0 for others; Black was coded 1 for Black faces and 0 for others. Angry unit activation was centered.

[†] $p < .10$.

* $p < .05$.

** $p < .01$.

*** $p < .001$.

Table 5

Predicting Impressions of Generalization Faces From Activation of the Happy Output Unit Controlling for Attractiveness, Babyfacedness, and Face Sex

Regression step	Predictor	Competence			Danger			Likeable			Expressive/loud		
		B	(SE)	β	B	(SE)	β	B	(SE)	β	B	(SE)	β
Step 1	Intercept	2.90	(.163)		6.04	(.193)		1.46	(.160)		3.31	(.184)	
	Korean (vs. White faces)	0.61	(.070)	.45***	-0.30	(.083)	-.19***	-0.14	(.069)	-.09*	-0.82	(.079)	-.50***
	Black (vs. White faces)	-0.21	(.069)	-.15**	0.37	(.082)	.22***	-.23	(.068)	-.15**	0.11	(.078)	.06
Step 2	Intercept	2.90	(.163)		6.03	(.191)		1.47	(.159)		3.32	(.184)	
	Korean (vs. White faces)	0.60	(.070)	.44***	-0.28	(.082)	-.18**	-0.16	(.068)	-.11*	-0.84	(.079)	-.51***
	Black (vs. White faces)	-0.25	(.082)	-.17**	0.50	(.096)	.30***	-0.33	(.080)	-.22***	0.01	(.092)	.01
Step 3	Happy unit activation	0.002	(.002)	.04	-0.01	(.002)	-.14**	0.004	(.002)	.11*	0.004	(.002)	.09 [†]
	Intercept	2.87	(.166)		6.06	(.195)		1.41	(.161)		3.30	(.187)	
	Korean (vs. White faces)	0.61	(.075)	.45***	-0.28	(.088)	-.18**	-0.10	(.073)	-.07	-0.80	(.084)	-.49***
	Black (vs. White faces)	-0.27	(.088)	-.19**	0.55	(.103)	.33***	-0.31	(.085)	-.20***	0.08	(.099)	.04
	Happy unit activation	-0.001	(.003)	-.02	-0.003	(.004)	-.07	-0.001	(.003)	-.01	0.002	(.004)	.06
	Happy × Korean	0.002	(.005)	.02	0.000	(.005)	-.003	0.01	(.004)	.12*	0.01	(.005)	.08
	Happy × Black	0.01	(.004)	.08	-0.01	(.005)	-.10	0.01	(.004)	.10	-0.002	(.005)	-.02

Note. Face sex was coded 0 for male faces and 1 for female faces. Korean was coded 1 for Korean faces and 0 for others; Black was coded 1 for Black faces and 0 for others. Happy unit activation was centered.

[†] $p < .10$.

* $p < .05$.

** $p < .01$.

*** $p < .001$.

Table 6

Predicting Impressions of Generalization Faces From Activation of the Surprise Output Unit Controlling for Attractiveness, Babyfacedness, and Face Sex

Regression step	Predictor	Competence			Danger			Likeable			Expressive/loud		
		B	(SE)	β	B	(SE)	β	B	(SE)	β	B	(SE)	β
Step 1	Intercept	2.90	(.163)		6.04	(.193)		1.46	(.160)		3.31	(.184)	
	Korean (vs. White faces)	0.61	(.070)	.45***	-0.30	(.083)	-.19***	-0.14	(.069)	-.09*	-0.82	(.079)	-.50***
	Black (vs. White faces)	-0.21	(.069)	-.15**	0.37	(.082)	.22***	-0.23	(.068)	-.15**	0.11	(.078)	.06
Step 2	Intercept	2.91	(.167)		5.99	(.197)		1.51	(.164)		3.25	(.188)	
	Korean (vs. White faces)	0.60	(.073)	.44***	-0.27	(.087)	-.17**	-0.16	(.072)	-.11*	-0.79	(.083)	-.48***
	Black (vs. White faces)	-0.24	(.092)	-.17*	0.45	(.109)	.27***	-0.30	(.091)	-.19**	0.21	(.104)	.12*
Step 3	Surprise unit activation	0.001	(.002)	.03	-0.002	(.002)	-.07	0.002	(.002)	.06	-0.003	(.002)	-.08
	Intercept	2.91	(.167)		6.00	(.198)		1.50	(.164)		3.24	(.189)	
	Korean (vs. White faces)	0.59	(.080)	.43***	-0.29	(.095)	-.19**	-0.15	(.078)	-.10 [†]	-0.77	(.090)	-.47***
	Black (vs. White faces)	-0.33	(.111)	-.23**	0.47	(.131)	.28***	-0.35	(.109)	-.22**	0.15	(.125)	.09
	Surprise unit activation	0.001	(.002)	.03	-0.001	(.003)	-.03	0.001	(.002)	.03	-0.004	(.003)	-.11
	Surprise × Korean	-0.003	(.004)	-.04	-0.002	(.004)	-.02	0.001	(.004)	.01	0.000	(.004)	.004
	Surprise × Black	0.004	(.004)	.08	-0.003	(.005)	-.05	0.004	(.004)	.07	0.01	(.005)	.07

Note. Face sex was coded 0 for male faces and 1 for female faces. Korean was coded 1 for Korean faces and 0 for others; Black was coded 1 for Black faces and 0 for others. Surprise unit activation was centered.

[†]
p < .10.

*
p < .05.

**
p < .01.

p < .001.