Social Cognition, Vol. 27, No. 6, 2009, pp. 813–833

EVALUATING FACES ON TRUSTWORTHINESS AFTER MINIMAL TIME EXPOSURE

Alexander Todorov, Manish Pakrashi, and Nikolaas N. Oosterhof
Princeton University

Previous studies have shown that trustworthiness judgments from facial appearance approximate general valence evaluation of faces (Oosterhof & Todorov, 2008) and are made after as little as 100 ms exposure to novel faces (Willis & Todorov, 2006). In Experiment 1, using better masking procedures and shorter exposures, we replicate the latter findings. In Experiment 2, we systematically manipulate the exposure to faces and show that a sigmoid function almost perfectly describes how judgments change as a function of time exposure. The agreement of these judgments with time-unconstrained judgments is above chance after 33 ms, improves with additional exposure, and does not improve with exposures longer than 167 ms. In Experiment 3, using a priming paradigm, we show that effects of face trustworthiness are detectable even when the faces are presented below the threshold of objective awareness as measured by a forced choice recognition test of the primes. The findings suggest that people automatically make valence/trustworthiness judgments from facial appearance.

Person impressions are often formed rapidly and spontaneously from minimal information (Todorov & Uleman, 2003; Uleman, Blader, & Todorov, 2005). One rich source of such information is facial appearance and there is abundant research about the effects of facial appearance on social outcomes (e.g., Blair, Judd, & Chapleau, 2004; Eberhardt, Davies, Purdie-Vaughns, & Johnson, 2006; Hamermesh & Biddle, 1994; Hassin & Trope, 2000; Langlois et al., 2000; Montepare & Zebrowitz, 1998; Zebrowitz, 1999). For example, inferences of competence, based solely on facial appearance, predict the outcomes of the U.S. congressional (Todorov, Mandisodza, Goren, & Hall, 2005) and gubernatorial elections (Ballew & Todorov, 2007; Hall, Goren, Chaiken, & Todorov, 2009), and inferences of dominance predict military rank attainment (Mazur, Mazur, & Keating, 1984; Mueller & Mazur, 1996).

We thank Amir Goren, Chris Olivola, Chris Said, Crystal Hall, Sara Verosky, and Sean Baron for their comments on previous drafts of this manuscript. This research was supported by National Science Foundation Grant BCS-0446846.

Correspondence concerning this article should be addressed to Alexander Todorov, Department of Psychology, Princeton University, Princeton, NJ 08540; E-mail: atodorov@princeton.edu.

Although it has been argued that trait inferences from facial appearance are rapid and automatic (e.g., Todorov et al., 2005), there have been only a few studies showing that such inferences are made after minimal time exposure to faces (Ballew & Todorov, 2007; Bar, Neta, & Linz, 2006; Willis & Todorov, 2006). In two of the studies, the minimum exposure was 100 ms (Ballew & Todorov, 2007; Willis & Todorov, 2006). Bar et al. (2006) used shorter exposures but failed to find effects on judgments after subliminal presentation of faces. Following this work, the first objective of the current studies was to identify the lower bound of time exposure after which people start discriminating between trustworthy- and untrustworthy-looking faces and to systematically map how explicit trustworthiness judgments change as a function of time exposure to faces. The second objective was to test whether effects of face trustworthiness are detectable even when the faces are presented below the threshold of objective awareness.

We focus on judgments of trustworthiness because they are an excellent approximation of the general valence evaluation of faces (Oosterhof & Todorov, 2008). Further, these judgments are made after as little as 100 ms exposure to a face and are faster than a variety of other personality judgments (Willis & Todorov, 2006).

Although people make multiple trait judgments from facial appearance, these judgments are highly correlated with each other. To a large extent, these inter-trait correlations reflect the valence evaluation that permeates social judgments (Kim & Rosenberg, 1980; Rosenberg, Nelson, & Vivekananthan, 1968; cf. Osgood, Suci, & Tennenbaum, 1957). Recently, in a series of behavioral studies, to obtain a measure of the valence evaluation of faces, we first identified trait dimensions that people spontaneously use to characterize faces (Oosterhof & Todorov, 2008). Then, emotionally neutral faces were rated on these dimensions and the mean trait judgments were submitted to a principal components analysis to identify the underlying structure of these judgments. The first principal component accounted for 63% of the variance of these judgments and reflected the evaluative meaning of the trait dimensions. All positive judgments (e.g., trustworthy, intelligent) had positive loadings and all negative judgments (e.g., weird, mean) had negative loadings on this component. Trustworthiness judgments had the highest loading (.94) on the evaluation component and were practically uncorrelated with the second component (-.06), which accounted for 18% of the variance and could be interpreted as dominance evaluation. This finding was robust with respect to the face stimuli and the set of judgments used to estimate the principal components (see also Todorov, Said, Engell, & Oosterhof, 2008). These findings suggest that trustworthiness judgments best approximate the valence dimension of face evaluation. In other words, a single trustworthiness judgment can serve as a reliable proxy of the valence evaluation of novel faces.

The behavioral findings that trustworthiness judgments reflect the valence evaluation of faces are consistent with studies suggesting that the amygdala, a subcortical brain structure critical for fear conditioning (LeDoux, 2000; Phelps & LeDoux, 2005) and detection of threats in the environment (Amaral, 2002; Davis & Whalen, 2001; Whalen, 1998), plays a key role in such judgments (Todorov & Engell, 2008). For example, Adolphs, Tranel, and Damasio (1998) showed that patients with bilateral amygdala damage are impaired in their discrimination of trustworthy- and untrustworthy-looking faces. Consistent with the Adolphs et al. (1998) findings, functional neuroimaging studies show that untrustworthy-looking faces evoke stronger activity in the amygdala than trustworthy-looking faces (Engell, Haxby,

& Todorov, 2007; Todorov, Baron, & Oosterhof, 2008; Winston, Strange, O'Doherty, & Dolan, 2002).

Thus, convergent evidence from behavioral and functional neuroimaging studies suggest that judgments of trustworthiness from faces (a) approximate valence evaluation of faces; (b) are rapidly and spontaneously computed; and (c) have an identifiable neural signature. In this article, we further explore the minimal conditions under which these judgments are made.

EXPERIMENT 1

Experiment 1 was modeled after the studies reported in Willis and Todorov (2006). Willis and Todorov (2006) studied five trait judgments from facial appearance: attractiveness, likeability, trustworthiness, competence, and aggressiveness. Faces were presented for 100 ms, 500 ms, or 1000 ms and participants were asked to make one of the five judgments. Attractiveness was included as a benchmark against which to compare the other four judgments. In contrast to these judgments, attractiveness is a property of facial appearance and prior studies have shown that judgments of attractiveness can be made after extremely brief presentations of faces (Locher, Unger, Sociedade, & Wahl, 1993; Olson & Marshuetz, 2005). For all five traits, judgments made after 100 ms exposure to the faces closely agreed with control judgments made in the absence of time constraints. More importantly, this agreement did not improve with additional time exposure, suggesting that 100 ms exposure is sufficient for people to form a person impression.

The agreement for judgments of trustworthiness was as high as the agreement for judgments of attractiveness. The response times for these two judgments were also almost identical and faster than the response times for the judgments of competence, likeability, and aggressiveness, further demonstrating the efficiency of trustworthiness judgments. However, this study did not identify the lower bound of time presentation sufficient for these judgments.

The objective of the current experiment was to replicate Willis and Todorov's findings using 50 ms exposure to faces. In addition, this experiment used a larger sample size and better experimental procedures than the Willis and Todorov studies. First, we collected new criterion judgments that were more closely related to the experimental judgments. Second, we used better masking procedures to ensure that the faces were presented for the intended exposure time.

METHOD

Participants. Eighty-three undergraduate students from Princeton University participated in the experiment for partial course credit.[1]

Face Stimuli and Criterion Judgments. Photographs were taken from the Karolinska Directed Emotional Faces set (Lundqvist, Flykt, & Öhman, 1998). These were

1. Four participants in Experiment 1 provided the same response to all trials and were replaced by new participants. This was the case for four participants in Experiment 2. In both cases, including the data for these participants did not change the results.

photographs of amateur actors between 20 and 30 years of age with no beards, mustaches, earrings, eyeglasses, or visible make-up, all wearing gray T-shirts. We used frontal head-shot photographs of 33 males and 33 females with neutral expressions and a direct gaze. The same stimuli were used by Willis and Todorov (2006).

However, we used new criterion judgments of trustworthiness. To obtain criterion judgments, Willis and Todorov (2006) had participants (n = 42 to 43) rate a subset of the faces on multiple traits and the faces were rated as grayscale images in a questionnaire format. These procedures introduce additional (and theoretically uninteresting) differences with the procedures used in the experiments and contribute to reduced correlations between experimental and criterion judgments. In a study described in detail elsewhere (Engell, Haxby, & Todorov, 2007, Study 1), we used a larger sample size (n = 129) and each participant rated all 66 faces. Each face was rated only on trustworthiness and the face images were presented in color the way they were presented in the experiments. The faces were rated on a 9-point scale ranging from 1 (Not at all trustworthy) to 9 (Extremely trustworthy). The ratings (M = 4.63; SD = 0.96) were highly reliable (Cronbach's alpha = 0.98). These ratings were used as the criterion judgments in Experiments 1 and 2.

The criterion judgments served as standards against which to compare the experimentally obtained judgments, with higher correlations between the two indicating greater correspondence between judgments obtained under limited time exposure and judgments obtained without time constraints.

Procedures. The procedures were modeled after the procedures of the experiments reported in Willis and Todorov (2006). The major differences were that (a) faces were presented for 50, 100, or 500 ms and (b) each face presentation was followed by a perceptual mask. Participants were told that this was a study about first impressions and that they should make their decisions as quickly as possible. The instructions emphasized that photographs would be presented for very brief periods of time and that we, the experimenters, were primarily interested in their first impression or gut reaction. The experiment started with 2 practice trials in order to familiarize participants with the task.

For the experimental trials, the 66 faces were randomly divided into 3 sets of 22 such that each group had the same number of male and female faces. Using these 3 sets of faces, we created 3 experimental versions by counterbalancing the sets with the exposure time (50, 100, or 500 ms). For example, each face from the first set of faces was presented for 50 ms in version 1, for 100 ms in version 2, and for 500 ms in version 3. Participants were randomly assigned to one of the 3 experimental versions. For each participant, 22 of the faces were presented for 50 ms, 22 for 100 ms, and 22 for 500 ms. Because we were interested in first impressions, each face was presented only once. Thus, the total number of trials was 66. The order of trials was randomized for each participant by the computer, that is, the levels of exposure time were randomly intermixed.

Each trial started with a fixation point (+) presented for 500 ms at the center of the screen. Then a face was presented either for 50 ms, 100 ms, or 500 ms and immediately replaced by a mask. Both the face images and the mask were 295 x 400 pixels bitmaps, and the approximate distance of the participant from the computer screen was 65 cm. The mask was composed of small facial segments of various faces, which were rearranged to form a jumbled, mosaic image (Figure 1a).

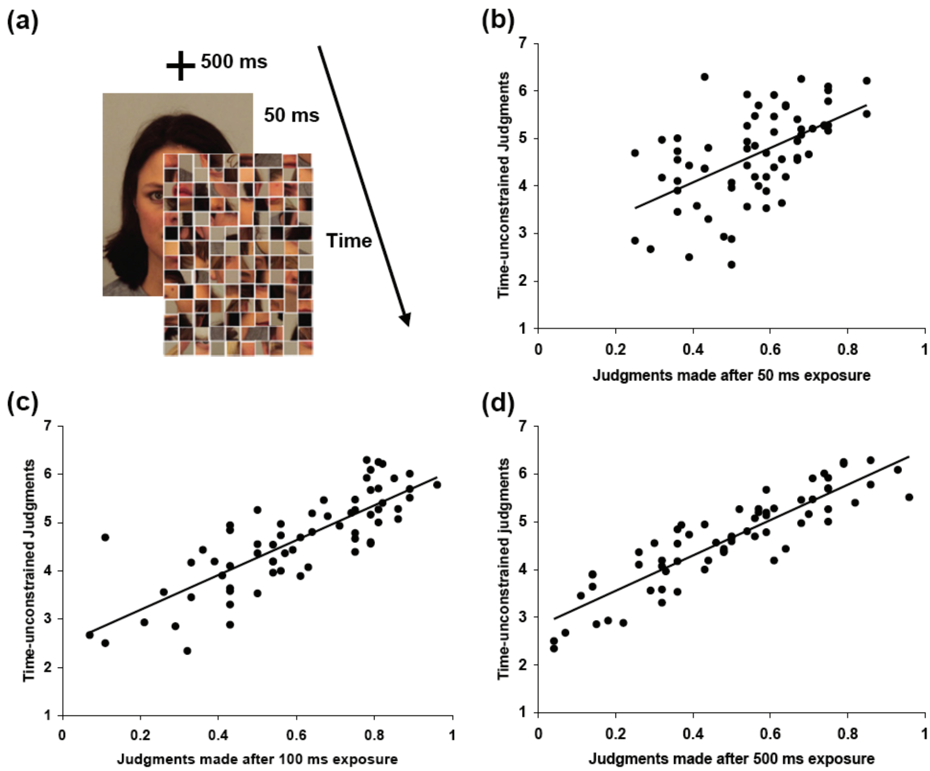

Figure 1. Example of an experimental trial and scatter plots of judgments of trustworthiness made after limited time exposure to faces and criterion judgments made in the absence of time constraints (Exp. 1). Each face was preceded by a fixation point and immediately masked after its presentation (a). Participants made a binary "yes/no" trustworthiness decision about the face. Scatter plots of judgments made in the absence of time constraints and judgments made after 50 ms exposure (b); judgments made after 100 ms exposure (c); and judgments made after 500 ms exposure (d). Each point represents a face. The line represents the best fitting line. The judgments made in the absence of time constraints – plotted on the y-axis – were made on a 9-point scale ranging from 1 (not at all trustworthy) to 9 (extremely trustworthy). The judgments made after limited time exposure – plotted on the x-axis – were binary yes/no judgments and the mean trustworthiness is a proportion.

A question "Is this person trustworthy?" was presented above the mask. The mask remained on screen until the participant's response. Participants responded by pressing the "/" (slash) key on the keyboard which was labeled "yes" or by pressing the "Z" key on the keyboard which was labeled "no." Following the "yes/no" trustworthiness judgment, the next screen asked participants to rate how confident they were in their judgment. This judgment was made on a 7-point scale, ranging from 1 (least confident) to 7 (most confident). Participants responded by using the number keys at the top of the keyboard. The inter-trial interval was 1500 ms.

Preliminary Analyses. In Experiments 1 and 2, some of the analyses used correlations as dependent variables. For example, in Experiment 1, for each participant we computed the correlation between their judgments of the faces made after 50 ms exposure and the mean criterion judgments. Then, we compared these correlations with the corresponding correlations for 100 and 500 ms exposures. For all analyses using correlations as dependent variables, we transformed the raw correlations into Fisher z-scores for statistical analyses. This is a standard procedure normalizing the distribution of correlations.

RESULTS

Trustworthiness Judgments. We conducted two analyses to test how time exposure affects judgments. In the first analysis, we correlated the mean judgments of the experimental participants for each time exposure with the mean criterion judgments. As shown in Figure 1b, even after 50 ms exposure, judgments correlated with the criterion judgments, $r(66) = .55$, $p < .001$. This correlation increased to .81 for 100 ms (Figure 1c) and .89 for 500 ms exposure (Figure 1d). The Williams's test for dependent correlations (Steiger, 1980) showed that both the difference between the correlations for 50 ms and 100 ms, $t(65) = 4.47$, $p < .001$, and the difference between the correlations for 100 and 500 ms were significant, $t(65) = 2.31$, $p < .02$. However, the effect size was twice as large for the increase in the correlation from 50 to 100 ms (Pearson $r = .48$) exposure as that for the increase from 100 to 500 ms exposure ($r = .28$).

Aggregating judgments across participants produces higher correlations than analysis at the level of individuals and it is important to show that these analyses are consistent. Thus, the second analysis was conducted at the level of the individual participants. For each participant, we computed the point-biserial correlation between their yes/no trustworthiness judgments of the 66 faces and the mean criterion judgments for each of the three levels of time exposure. The average correlation was .18 ($SD = .23$) after 50 ms exposure, .36 ($SD = .21$) after 100 ms exposure, and .43 ($SD = .17$) after 500 ms exposure. The average correlation for the 50 ms exposure was significantly higher than zero, $t(82) = 7.10$, $p < .001$. Both the difference between the correlations at 50 and 100 ms, $t(82) = 5.97$, $p < .001$, and at 100 and 500 ms, $t(82) = 2.26$, $p < .026$, were significant. As in the case of the aggregated judgments, the effect size was larger for the change from 50 to 100 ms ($r = .55$) than for the change from 100 to 500 ms ($r = .24$). Consistent with this difference, a test of the quadratic trend for the change in correlations was significant, $F(1, 82) = 4.72$, $p < .033$.

Confidence in Judgments. Confidence in judgments increased with the increase in time exposure to faces, $F(2, 164) = 111.39$, $p < .001$. Confidence increased from 2.79 ($SD = 1.04$) at 50 ms to 3.47 ($SD = 1.02$) at 100 ms, $t(82) = 11.80$, $p < .001$, and to 3.94 ($SD = 0.95$) at 500 ms, $t(82) = 6.46$, $p < .001$ (for the difference between 100 and 500 ms).

DISCUSSION

Using better procedures—perceptual masking of faces, more reliable criterion judgments, and a larger sample size—we replicated the findings of Willis and Todorov (2006) and identified another lower bound of trustworthiness judgments. Specifically, people were able to make these judgments even after 50 ms exposure to faces. The change in judgments and confidence in judgments were much more pronounced in the increase in exposure from 50 to 100 ms than in the increase from 100 to 500 ms. These findings suggest that most of the information needed to make a trait judgment is obtained within the first 100 ms exposure to the face.

EXPERIMENT 2

In the first experiment, participants were presented with each face only once. Although this procedure captures the idea of testing the formation of first impressions best, it effectively reduces the number of experimental trials and makes it difficult to systematically explore how judgments change as a function of exposure time to faces. Given that the first experiment established that people are able to make trustworthiness judgments after minimal time exposure on a limited number of trials, we decided to present the faces multiple times at different time exposures. In order to model quantitatively how judgments change as a function of time exposure, we used eight different exposures (17, 33, 50, 67, 100, 167, 500 ms, and unlimited time). Consistent with prior findings (Bar et al., 2006), we expected that (a) participants would be unable to discriminate between trustworthy- and untrustworthy-looking faces at extremely short exposures (e.g., 17 ms); (b) with the increase in time exposure from 33 to 100 ms, there would be a relatively large increase in the agreement of judgments made after limited exposure with criterion judgments; and (c) there would be relatively little change in this agreement with longer exposures.

As was described in the introduction, we recently showed that judgments of trustworthiness are highly correlated with the valence evaluation of faces (Oosterhof & Todorov, 2008). Because these analyses were conducted on the same set of faces as the one used in the current experiment, it was possible to test whether trustworthiness judgments made after limited time exposure to faces reflect general valence evaluation. For this analysis, we used the face values on the first (valence evaluation) component estimated from all trait judgments used by Oosterhof & Todorov except trustworthiness to avoid a biased solution with respect to trustworthiness. The correlation between this component—a linear combination of 12 trait judgments—and the mean criterion trustworthiness judgments was .93, $p < .001$. In contrast, the correlation between the second (dominance evaluation) component and the mean criterion judgments was -.13, $p = .29$.

Even if trustworthiness judgments made after rapid exposure to faces primarily reflect valence evaluation of faces, it is possible that these judgments also have components specific to trustworthiness evaluation. To test this hypothesis, we correlated the judgments made after rapid exposure with the mean criterion judgments after removing the shared variance of the latter with the valence component. If trustworthiness judgments made after rapid exposure reflect not only valence but also trustworthiness specific evaluations, these correlations should be significantly higher than zero.

METHOD

Participants. Forty-three undergraduate students from Princeton University participated in the study for partial course credit.

Procedures. Participants were told that the study is about impressions of trustworthiness from faces and that the faces would be presented for varying amounts of time. The instructions stressed that some of the pictures would be presented for an extremely short amount of time and that participants should use their "gut instinct" to respond. The 66 faces were presented in 8 blocks of incremental exposure times: 17, 33, 50, 67, 100, 167, 500 ms, and unlimited exposure. Within each block, the order of the trials was randomized for each participant. The total number of trials was 528 (66 faces x 8 blocks).

Each trial started with a 500 ms fixation point presented at the center of the screen. The fixation was followed by a face presented for its respective exposure time and replaced by a mask (Figure 1a). As in Experiment 1, the images and mask were 295 x 400 pixels bitmaps. A question "Is the person you just saw trustworthy?" was presented above the mask and two response choices ("Yes" and "No") were presented below the mask. Participants responded by pressing either the "Z" key on the keyboard, which was labeled "YES", or the "/" key on the keyboard, which was labeled "NO." In the final block (no time constraints), each face was presented until the participant responded. The inter-trial interval was 1000 ms.

RESULTS

As shown in Fig. 2a, the correlation between the mean judgments for each exposure and the mean criterion judgments systematically increased as function of the exposure time. At 17 ms exposure, the correlation was not significantly different from zero. The correlation increased substantially with the increase in exposure from 33 to 100 ms (from .22 to .72) and did not increase after 167 ms. A sigmoid function ($r = e^{0.17 - 55.69/\mathrm{time}}$) accounted for 96.6% of the variance of these data. To fit this function, we could not use the self-pace presentation condition because the presentation time varied across participants. However, the correlation between the criterion judgments and the judgments in this condition (.75) was not higher than the correlations for the 167 ms condition (.79) and the 500 ms condition (.79).

We conducted the same analysis at the level of the correlations for individual participants. Although the correlations were lower, the pattern of results was identical (Figure 2b). A sigmoid function ($r = e^{-0.87 - 59.29/\mathrm{time}}$) accounted for 95.4% of the variance. At 17 ms exposure, the average correlation was not significantly different from zero, $t < 1$, but at 33 ms it was significantly higher than zero, $t(42) = 2.17$, $p < .036$. The difference between the correlations at 167 ms and 500 ms conditions was not significant, $t < 1$. Similarly, the difference between the correlations at 500 ms and the self-pace conditions was not significant, $t < 1$.

The correlation between the valence evaluation component, estimated from 12 trait judgments, and the judgments made after rapid exposure closely approximated the correlation between these judgments and the mean criterion trustworthiness judgments (Figure 2). The sigmoid functions describing how these correlations (for criterion trustworthiness judgments and valence evaluation) changed as a function of exposure time were perfectly correlated. This was the case for both aggregated and individual level correlations.

Nevertheless, judgments made after rapid exposure were positively correlated with the residuals of the mean criterion judgments after regressing the latter on the valence component (Figure 2). For the aggregated judgments, these were sig-

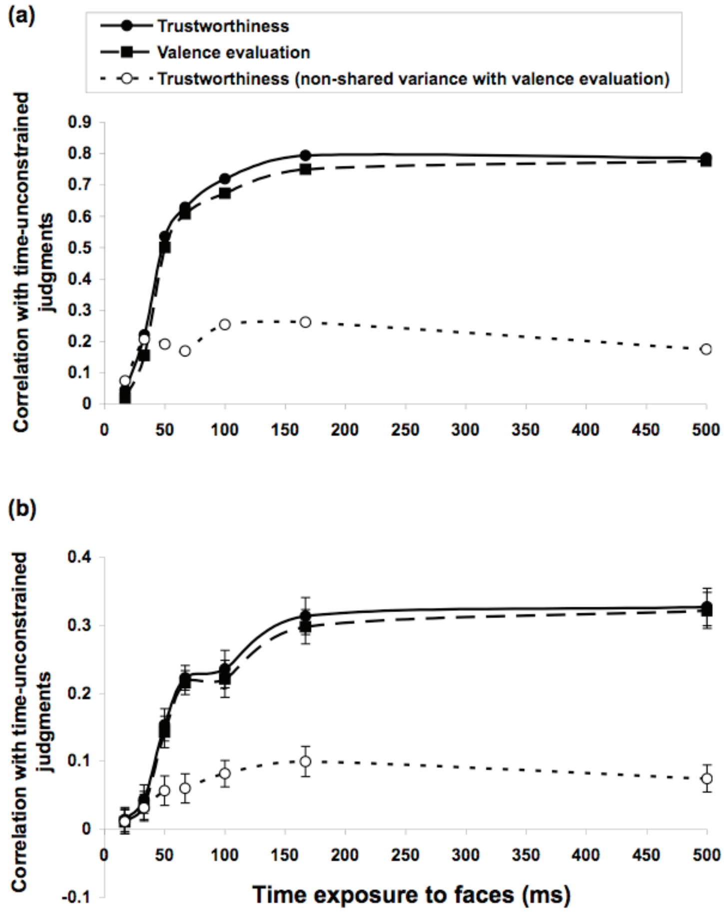

FIGURE 2. Correlation of judgments of trustworthiness made after limited time exposure to faces with criterion judgments of trustworthiness made in the absence of time constraints, valence evaluation of faces estimated from a principal components analysis of trait judgments (Oosterhof & Todorov, 2008), and residual variance of criterion trustworthiness judgments after removing shared variance with valence evaluation (Exp. 2). Faces were presented for 17, 33, 50, 67, 100, 167, and 500 ms. Correlations at the level of mean trustworthiness judgment averaged across participants (a). Correlations at the level of individual participants (b). The means represent the average of the correlations between individual judgments and the criterion judgments made in the absence of time constraints. Error bars show the standard error of the mean. The valence evaluation is a linear combination of 12 trait judgments other than trustworthiness (see for details Oosterhof & Todorov, 2008).

nificantly higher than zero for 100 and 167 ms exposures. For the individual level data, the correlations were significantly higher than zero for 50 ms and longer exposures, *t*s(42) > 2.61, *p*s < .012.

DISCUSSION

Using multiple presentation times, we showed that judgments of trustworthiness changed systematically as a function of the exposure time to faces. After presentation time (17 ms) possibly precluding conscious perception of the faces, participants were unable to discriminate trustworthy- from untrustworthy-looking faces. However, after 33 ms exposure, participants discriminated between these faces.[2] The correlation with the mean criterion judgments made in the absence of time constraints increased dramatically with the increase in exposure from 33 to 100 ms and relatively little with the increase in exposure from 100 to 167 ms. Additional increases in time exposure did not improve this correlation. The pattern was the same for the correlation with the general valence evaluation component estimated from the data of Oosterhof & Todorov (2008). These findings suggest that trustworthiness judgments made after brief exposures to faces reflect general valence evaluation of faces. Interestingly, although the correlations with valence evaluation closely mirrored the correlations with the criterion judgments, judgments made after brief exposures also correlated with the variance of the criterion judgments that was not shared with the valence component. This finding suggests that trustworthiness judgments made after brief exposures also reflect information specific to trustworthiness.

EXPERIMENT 3

The findings of Experiment 2 are consistent with the findings of Bar et al. (2006) who reported that threat judgments made after 39 ms (but not judgments made after 26 ms) correlated highly with judgments made after 1700 ms. These findings suggest that trait judgments from faces are not made after subliminal exposure. However, there are two potential problems with this conclusion. First, explicit trait judgments may not be sensitive enough to detect differences in perception of trustworthy and untrustworthy faces, or threatening and nonthreatening faces,[3] after subliminal exposures. Second, faces of actors that are typically used in psychology experiments may not provide a sufficient range of differences on the trait dimension.

In this experiment, we used a priming task to test whether differences in perception of trustworthy and untrustworthy faces can be detected after subliminal presentation. To provide a large range of differences on trustworthiness, we used faces generated by a validated computer model of face trustworthiness (Oosterhof & Todorov, 2008). Using this model, it is possible to generate an unlimited number of faces and manipulate their trustworthiness. For the current experiment, we used three levels of trustworthiness (-3, 0, +3 *SD* on a normally distributed dimension; see Fig. 3 for examples). In this range, faces are perceived as emotionally neutral, but there are large differences in perception of trustworthiness.

On the critical trials, untrustworthy (-3 *SD*) or trustworthy faces (3 *SD*) were presented for 20 ms and immediately masked by a neutral face (0 *SD*) with respect to trustworthiness. The participant's task was to judge this neutral face on trust-

---

2. We replicated this finding in a separate experiment, in which 41 participants judged the 16 least trustworthy and the 16 most trustworthy faces. The faces were presented once for 33 ms. Participants were more likely to judge the trustworthy faces ($M = .52$, $SD = .23$) as more trustworthy than the untrustworthy faces ($M = .46$, $SD = .22$), $t(40) = 2.56$, $p < .014$ (effect size $r = .38$).

3. As shown by Oosterhof and Todorov (2008), threatening facial features can be obtained from a linear combination of untrustworthy and dominant facial features.

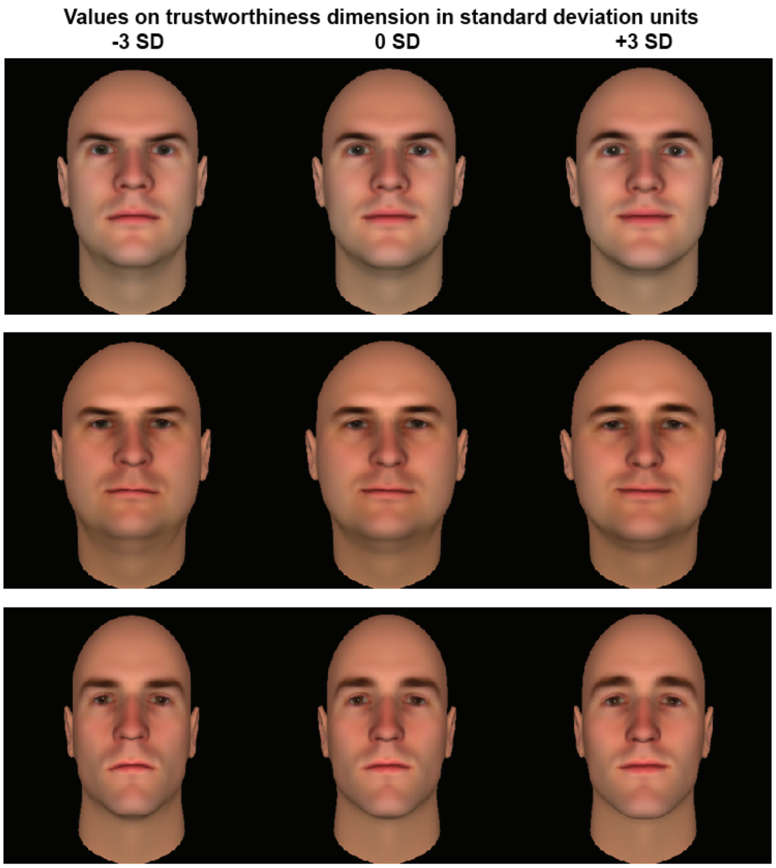

FIGURE 3. Examples of untrustworthy, neutral, and trustworthy versions of faces generated by a computer model of face trustworthiness (Oosterhof & Todorov, 2008). Each rows shows a different facial identity. The first column shows the untrustworthy version of the faces (-3 SD on the trustworthiness dimension), the second column shows the neutral version (0 SD), and the third column shows the trustworthy version (+3SD).

worthiness (Figure 4a). If face trustworthiness is evaluated even after presentations precluding conscious awareness of the faces, participants should judge neutral faces as more untrustworthy when these faces are preceded by untrustworthy than by trustworthy faces.

To ensure that participants were not aware of the primes, we included an objective test of awareness at the end of the experiment (Cheesman & Merikle, 1984). After the priming task, participants were informed about the presence of primes and asked to perform a forced choice recognition task. Specifically, the priming trials were identical to the trials in the first part of the experiment. However, after each priming trial, participants were presented with two faces and asked to guess which face preceded the target (neutral) face (Figure 4b).

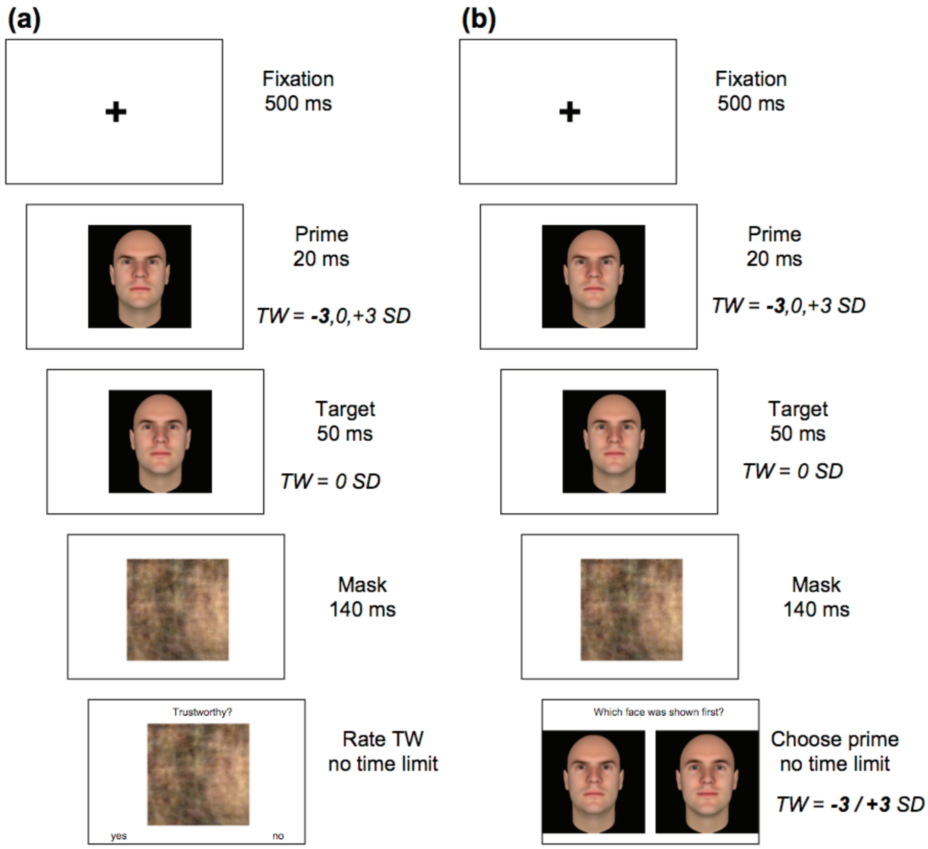

FIGURE 4. Examples of experimental trials in the subliminal priming paradigm (Exp. 3). Priming trial from the first part of the experiment (a). Recognition trial from the second part of the experiment (b). Trustworthy and untrustworthy faces (TW = +3 and -3 SD, respectively) were presented for 20 ms and immediately masked by a neutral face (TW = 0 SD). The examples show an untrustworthy face prime. The trustworthy equivalent can be seen in the lower right part of the figure (recognition trial).

METHOD

Participants. Thirty-two undergraduate students from Princeton University participated in the experiment for partial course credit.

Face Stimuli. We used the Facegen Modeller program (http://facegen.com) version 3.1 to generate faces and the trustworthiness face model developed by Oosterhof and Todorov (2008). Facegen creates 3D faces whose shape and texture can be adjusted on multiple dimensions. The face model of Facegen (Blanz & Vetter, 1999; Singular Inversions, 2006) is based on a database of male and female faces that were laser-scanned in 3D. Using a principal component analysis, a model was constructed so that each face can be represented by a limited number of independent components. Oosterhof and Todorov worked with the 50 components, representing symmetric facial shape, to build a model of face trustworthiness. Specifically, they used trustworthiness judgments of randomly generated faces to build a dimension in the 50-dimensional space optimal in changing face trustworthiness. Subsequent studies showed that judgments of faces generated by the trustworthiness model tracked the trustworthiness predicted by the model.

For the current experiment, 150 Caucasian male faces were generated using Facegen. The faces were generated randomly with four constraints. First, Facegen's race controls were set to European to avoid judgments affected by ethnic stereotypes because a randomly generated face can be of any race. Second, facial attractiveness was increased to make the faces more similar to the photo-fitted real faces used in Experiments 1 and 2. Third, all faces were male because these faces look more natural without hair (i.e., bold). Fourth, although all faces were set to be emotionally neutral by default, we also set the mouth shape control, which moves the corners of the mouth up and down, to neutral to further ensure that the expressions are neutral.

Untrustworthy, neutral, and trustworthy versions were then created for each of the 150 face identities resulting in 450 faces that were used as stimuli in the experiment (see Fig. 3 for examples). Trustworthiness was manipulated using a custom created and validated Facegen control that changes facial features that are specific for trustworthiness (Oosterhof & Todorov, 2008). The untrustworthy, neutral, and trustworthy versions were produced by setting this trustworthiness control to -3, 0, and +3 *SD* respectively and saved as 400 x 400 pixels bitmaps. Previous validation studies of the computer model of face trustworthiness (Oosterhof & Todorov, 2008) have shown that people easily discriminate between these levels of trustworthiness. For example, the mean trustworthiness judgments for -3, 0, and +3 *SD* faces were 3.99 (*SD* = .74), 5.23 (*SD* = .30), and 5.95 (*SD* = .62), respectively.

Procedures. Participants were told that this was a study in two parts about first impressions and that they should rely on their "gut" instinct in judging the faces. Participants were presented with three practice trials in the beginning of each task to become familiar with the procedures. The experiment was run using Eprime on standard PCs.

The experiment was split into two sections. For the first part of the experiment, participants were asked to judge the trustworthiness of faces. Participants were presented with 150 trials: 100 priming and 50 baseline trials. For 100 of the trials, the target face was preceded either by an untrustworthy face prime or a trustworthy face prime. Participants were not informed about the presence of a prime stimulus. Each of the trials began with a fixation point (+) presented for 500 ms at the center of the screen. For the priming trials, an untrustworthy or trustworthy prime stimulus was presented for 20 ms, followed by a neutral target stimulus (with the same facial identity as the prime) presented for 50 ms. To ensure accurate stimulus presentation times of the primes, for each trial the images were preloaded while the fixation cross was displayed. The accuracy of stimulus presentation times was verified using a photodiode and a USB-1208FS data acquisition device (Measurement Computing, Norton, MA) sampling at 2kHz. Stimuli were presented on 17 inch CRT monitors with a 100Hz refresh rate at a viewing distance of approximately 65 cm. Faces were surrounded by a black square of 15.2 × 15.2 cm. The size of the faces was 11.6 cm vertically (top to chin) and 7.8 cm horizontally.

The position of the prime was offset to the target and mask randomly with ±5 pixels in both the horizontal and vertical direction, in order to avoid apparent facial movement that can affect judgments (Philips et al., 2004; Wiens, 2006). After the 50 ms presentation of the target face, the face was masked (Figure 4a). The mask was first shown for about 140 ms (±17 ms). The variability in time presentation was due to differences in the time for loading the mask in Eprime, which did not affect the prime, target, and mask onsets because they were preloaded in memory using image caching. Subsequently, without delay and at the same position, the mask was shown with the question "Trustworthy?" presented above the mask. The mask remained on the screen until the participant responded by pressing the "A" key on the keyboard which was labeled "yes" or by pressing the ";" (semicolon) key on the keyboard which was labeled "no." For the 50 baseline trials, the neutral target was presented for 70 ms and masked as in the priming trials. The inter-trial interval was 1000 ms. During the 150 trials each facial identity was shown once. The trustworthiness of the prime stimuli was counterbalanced in regards to facial identity between participants, and the order of the trials was randomized for each participant.

To test whether participants were aware of the prime stimuli, the second part of the study was a forced choice recognition task. Participants were informed about the presence of the prime stimuli and then presented with the same 150 trials as in the first part of the study. However, after the first presentation of the mask (Figure 4b), two faces with high (3 *SD*) and low (-3 *SD*) trustworthiness originating from the same neutral face (0 *SD*) were shown. Above the faces was the question "Which face was shown?" and participants then selected the left or right stimulus by pressing the "Z" or "M" key on the keyboard, respectively. The position of the correct stimulus (left or right) was counterbalanced across participants and the order of the trials was randomized for each participant.

RESULTS

Manipulation Check of Awareness. A calibrated observer should be more likely to recognize trustworthy faces when such faces are primed (correct recognition) than when untrustworthy faces are primed (false recognition) and should be at chance (.50) on the baseline trials in which no primes were presented. As shown in Fig. 5a, participants were at chance discriminating between trustworthy and untrustworthy primes, $F < 1$, $p = .53$, for the overall effect, and $F < 1$, $p = .56$ for the difference between correct and false recognition rates for the priming trials. We also conducted a signal detection analysis of the recognition performance on the priming trials. The d-prime ($M = .03$, $SD = .25$) was not significantly different from zero, $t < 1$, $p = .58$

Priming Effects. As shown in Fig. 5b, participants were more likely to judge neutral faces as untrustworthy when they were preceded by untrustworthy than by trustworthy faces, $F(2, 62) = 6.42$, $p < .003$, for the overall effect, and $F(1, 31) = 10.90$, $p < .002$, for the linear trend. Subsequent analyses also showed that whereas the difference between judgments in the negative prime and the baseline no-prime conditions was significant, $t(31) = 2.45$, $p < .020$, the difference between judgments in the latter condition and the positive prime condition was not significant, $t(31) = 1.21$, $p = .24$, though it was in the expected direction (Figure 5b).

Although participants were at chance in their recognition performance, it is theoretically possible that recognition performance can account for the priming effects. For example, it could be that individuals with better than chance recogni-

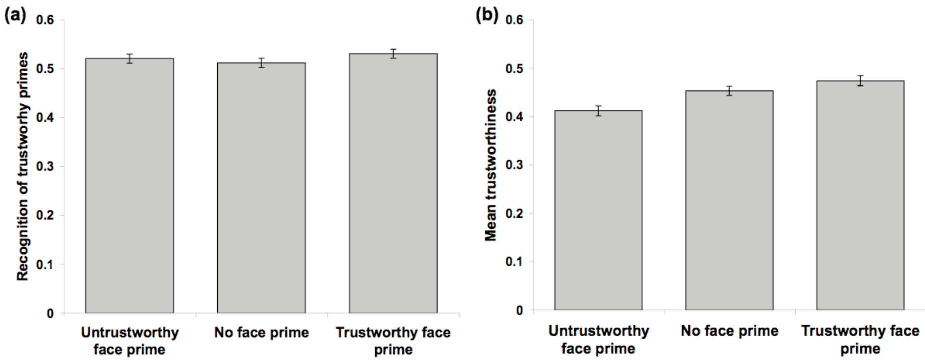

FIGURE 5. Recognition and priming performance (Exp. 3). Proportion of recognition of trustworthy primes (a). The first two bars show false recognition because either untrustworthy primes (first bar) or no primes (second bar) were presented. The third bar shows correct recognition. Trustworthiness judgments as a function of priming condition (b). Error bars show the within-subjects error of the mean.

tion performance show large priming effects whereas individuals with chance or worse than chance performance do not show any priming effects. However, at the individual level, the priming effect (measured as the difference between the judgments in the positive and negative prime conditions) and the recognition effect (measured as the difference between the correct and false recognition of positive primes) were practically uncorrelated (r = .09, p = .62). Similarly, the correlation between the priming effect and d-prime was .12, p = .52. We also binned participants into four groups in terms of their recognition performance, ranging from worst performance (d-prime < -.16) to best performance (d-prime > .25) and submitted the judgment to 3 (prime condition) X 4 (recognition group) mixed subjects ANOVA. The only significant effect in this analysis was the priming effect, $F(2, 56) = 6.40$, $p < .003$ ($F(1, 28)=10.84$, $p < .003$ for the linear effect); $F < 1$ for the interaction.

DISCUSSION

We showed that when the task is sensitive enough, it is possible to detect effects of face trustworthiness on judgments after subliminal presentation of faces. Although participants were at chance discriminating between presented primes and filler faces, they were more likely to judge neutral faces as more untrustworthy when these faces were primed with untrustworthy than with trustworthy faces. Additional analyses showed that the variation in priming performance couldn't be explained by the variation in recognition performance. This finding is consistent with prior demonstration in the literature that emotional faces can be perceived even if they are presented below the level of conscious awareness (e.g., Whalen et al., 1998; Whalen et al., 2004; Winkielman, Berridge, & Wilbarger, 2005). To the best of our knowledge, this is the first demonstration of a subliminal priming effect with emotionally neutral faces of the same race that vary on a social dimension.

GENERAL DISCUSSION

The objectives of the first two experiments were to identify the minimal time exposure to faces sufficient for people to make trustworthiness judgments and to systematically map how judgments change as a function of time exposure. The first experiment replicated Willis and Todorov's (2006) findings using a lower bound of 50 ms exposure, better masking procedures, better criterion judgments, and a larger sample size. In both experiments, with the increase in time exposure, trustworthiness judgments substantially improved. In Experiment 2, a sigmoid function almost perfectly described how judgments changed as a function of time exposure. After 17 ms exposure, participants' judgments were not significantly different from chance. With the increase in exposure from 33 to 100 ms, there was a large improvement in judgments. In contrast, with the increase from 100 to 167 ms, there was relatively little improvement. Judgments made after exposures longer than 167 ms did not change.

Although 33 ms exposure to faces was sufficient for trait judgments, these judgments were at chance after subliminal exposure to faces. This finding, coupled with the Bar et al. (2006) finding that participants were able to discriminate between faces that appear threatening and nonthreatening after 39 ms exposure but not after 26 ms, suggest that it is unlikely that people can make trait judgments after face exposures below their subjective awareness. One possible exception is judgments of attractiveness. For example, Olson and Marshuetz (2005), using forward and backward masking procedures, showed that participants were able to make attractiveness judgments after 17 ms exposure to faces although the face was not immediately masked after its presentation. Instead, the backward mask was presented after another 13 ms. Thus, the total exposure before the backward mask was presented was 26 ms.

Yet, all of the previous studies used an explicit judgment task in which the briefly presented face was judged on the dimension of interest. Such tasks may not be sensitive enough to detect effects of appearance on trait judgments after subliminal exposures. In Experiment 3, we used a priming paradigm in which trustworthy and untrustworthy faces presented for 20 ms were masked with a neutral face with respect to trustworthiness. Although participants were at chance discriminating between the primes in a forced choice recognition task—a measure of objective awareness—they judged the target-neutral faces as more untrustworthy when these faces were preceded by untrustworthy faces than when the faces were not preceded by a prime or preceded by trustworthy faces. These findings suggest that the trustworthiness of novel faces can be evaluated even when participants are unaware of the presence of the face, consistent with functional neuroimaging studies showing that the amygdala can be activated by the affective value of emotional faces of which participants are unaware (e.g., Pasley, Mayes, & Schultz, 2004; Whalen et al., 1998; Whalen et al., 2004; Williams, Morris, McGlone, Abbott, & Mattingley, 2004).

As mentioned in the introduction, functional neuroimaging studies show that untrustworthy-looking faces evoke a stronger amygdala response than trustworthy-looking faces (Engell et al., 2007; Todorov et al., 2008; Winston et al., 2002). Moreover, this effect is independent of the explicit task of the participants (Winston et al., 2002). For example, in Engell et al. (2007), participants ostensibly participated in a face memory task. They were presented with blocks of faces and asked to indicate whether a test face was in the block of preceding faces. Although this task did not demand person evaluation, the amygdala activation tracked the perceived trustworthiness of faces as measured by judgments obtained by an independent sample of participants. The more untrustworthy the face, the stronger was the amygdala response to the face. These findings, coupled with the findings of the current experiments, suggest that novel faces are automatically evaluated on valence/trustworthiness (Todorov & Engell, 2008).

A NOTE ON THE ACCURACY OF TRAIT JUDGMENTS FROM FACES

In our research, the measure of "goodness" of judgments is the extent to which they correspond to consensus judgments. This pertains to the reliability of these judgments but not to their validity. The fact that people agree that a person looks untrustworthy does not mean that the person is actually untrustworthy. Are trait judgments from faces accurate?

There is a large body of evidence that "thin slices" of behaviors can provide sufficient information for accurate social judgments (e.g., Albright, Kenny, & Malloy, 1988; Ambady, Hallahan, & Rosenthal, 1995; Ambady & Rosenthal, 1992; Borkenau & Liebler, 1992; Kenny, Horner, Kashy, & Chu, 1992). For example, judgments of extroversion from handshaking correlate with actual extroversion (Chaplin, Phillips, Brown, Clanton, & Stein, 2000), and personality judgments from personal websites correlate with self-reports on the big five dimensions of personality (Vazire & Gosling, 2004). However, nonverbal information in minimal interactions (e.g., Kenny et al., 1992) provides multiple cues about personality and what a person puts on his or her website is strategically selected to represent his or her personality to the world. A few of these cues are present in static images of faces, and judgments from these static images can be different from judgments from dynamic facial images (Rubenstein, 2005). Nevertheless, there have been studies finding significant correlations between various trait judgments from faces and measures of these traits, particularly extroversion (Borkenau & Liebler, 1992; Penton-Voak, Pound, Little, & Perrett, 2006). However, in most of these cases, the correlations are modest and judgments are not accurate for many other trait dimensions (e.g., agreeableness). There are also studies finding negative correlations between trait judgments and behaviors (Collins & Zebrowitz, 1995; Hassin & Trope, 2000; Zebrowitz, Andreoletti, Collins, Lee, & Blumenthal, 1998; Zebrowitz, Collins, & Dutta, 1998).

In the case of trustworthiness, there is evidence that judgments from still images of faces can have some degree of accuracy (Berry, 1990; Bond, Berry, & Omar, 1994). Berry (1990) found that judgments of honesty correlated with self-reports and judgments of acquaintances. Bond et al. (1994) found that judgments of honesty predicted willingness to participate in experiments that involved deceiving another participant. Specifically, participants who were rated as dishonest were more likely to express willingness to participate in these experiments. However, Zebrowitz, Voinescu, and Collins (1996) failed to find an overall correspondence between judgments of honesty from faces and clinical assessments of honesty. Additional analyses showed positive correlations for men who had a stable appearance of honesty across the life span, but negative correlations for women.

Although the evidence for the accuracy of trait judgments from faces is mixed, it would be puzzling, from an evolutionary point of view, to have an efficient system for making trait inferences from faces that does not deliver veridical inferences (Todorov, 2008). An alternative possibility—allowing for efficient but inaccurate trait judgments—is that these judgments are constructed from facial cues that have adaptive significance (Zebrowitz & Montepare, 2006, 2008). For example, according to the emotion overgeneralization hypothesis, resemblance of facial features to emotional expressions may be misattributed to personality dispositions (Knutson, 1996; Montepare & Dobish, 2003; Said, Sebe, & Todorov, 2009; Secord, 1958; Todorov, in press).

In particular, computer modeling of trustworthiness judgments from faces suggest that these judgments are grounded in similarity to emotional expressions that signal corresponding approach/avoidance behaviors to the perceiver (Oosterhof & Todorov, 2008; Todorov et al., 2008). Specifically, whereas increasing the trustworthiness of faces generated by a model of face trustworthiness resulted in happy expressions, decreasing their trustworthiness resulted in angry expressions. These findings suggest that subtle resemblance of neutral faces to happy expressions lead to attributions of trustworthiness and subtle resemblance to angry expressions lead to attributions of untrustworthiness, attributions that are not necessarily warranted.

CONCLUSIONS

We showed that trustworthiness judgments are made within a single glance of a face. As Macrae and colleagues have noticed, social cognition research on face perception has been exclusively focused on how categorical knowledge such as age and sex is activated from faces (Macrae, Quinn, Mason, & Quadflieg, 2005). The current research shows that personality inferences are also rapidly extracted from facial appearance. The paradigms introduced here can be easily extended to study how categorical and personality cues from faces are integrated in the process of person construal. For example, faces generated by the computer models for representing face variations on trait dimensions developed by Oosterhof & Todorov (2008) can also be manipulated on dimensions such as age, sex, and race and experiments manipulating the time exposure to faces can map how categorical and personality cues are extracted and integrated in person construal.

REFERENCES

Adolphs, R., Tranel, D., & Damasio, A. R. (1998). The human amygdala in social judgment. *Nature, 393*, 470-474.

Albright, L., Kenny, D. A., & Malloy, T. E. (1988). Consensus in personality judgments at zero acquaintance. *Journal of Personality and Social Psychology, 55*, 387-395.

Amaral, D. G. (2002). The primate amygdala and the neurobiology of social behavior: Implications for understanding social anxiety. *Biological Psychiatry, 51*, 11-17.

Ambady, N., Hallahan, M., & Rosenthal, R. (1995). On judging and being judged accurately in zero-acquaintance situations.

Journal of Personality and Social Psychology, 69, 518-529.

Ambady, N., & Rosenthal, R. (1992). Thin slices of expressive behavior as predictors of interpersonal consequences: A meta-analysis. *Psychological Bulletin, 111*, 256-274.

Ballew, C. C., & Todorov, A. (2007). Predicting political elections from rapid and unreflective face judgments. *Proceedings of the National Academy of Sciences of the USA, 104*(46), 17948-17953.

Bar, M., Neta, M., & Linz, H. (2006). Very first impressions. *Emotion, 6*, 269-278.

Berry, D. S. (1990). Taking people at face value: Evidence for the kernel of truth hypothesis. *Social Cognition, 8*, 343-361.

Blair, I. V., Judd, C. M., & Chapleau, K. M. (2004). The influence of Afrocentric facial features in criminal sentencing. *Psychological Science, 15*, 674-679.

Blanz, V., & Vetter, T. (1999). A morphable model for the synthesis of 3D faces. In *Proceedings of the 26th annual conference on computer graphics and interactive techniques*, 187-194.

Bond, C. F., Berry, D. S., & Omar, A. (1994). The kernel of truth in judgments of deceptiveness. Basic and Applied Social Psychology, 15, 523-534.

Borkenau, P., & Liebler, A. (1992). Trait inferences: Sources of validity at zero acquaintance. *Journal of Personality and Social Psychology, 65*, 645-657.

Chaplin, W. F., Phillips, J. B., Brown, J. D., Clanton, N. R., & Stein, J. L. (2000). Handshaking, gender, personality, and first impressions. *Journal of Personality and Social Psychology, 79*, 110-117.

Cheesman, J., & Merikle, P. M. (1984). Priming with and without awareness. *Perception & Psychophysics, 36*, 387-395.

Collins, M. A., & Zebrowitz, L. A. (1995). The contributions of appearance to occupational outcomes in civilian and military settings. *Journal of Applied Social Psychology, 25*, 129-163.

Davis, M., & Whalen, P. J. (2001). The amygdala: Vigilance and emotion. *Molecular Psychiatry, 6*, 13-34.

Eberhardt, J. L., Davies, P. G., Purdie-Vaughns, V. J., & Johnson, S. L. (2006). Looking deathworthy: Perceived stereotypicality of Black defendants predicts capital-sentencing outcomes. *Psychological Science, 17*, 383-386.

Engell, A. D., Haxby, J. V., & Todorov, A. (2007). Implicit trustworthiness decisions: Automatic coding of face properties in human amygdala. *Journal of Cognitive Neuroscience, 19*, 1508-1519.

Gosling, S. D., Rentfrow, P. J., & Swann, W. B., Jr. (2003). A very brief measure of the Big Five personality domains. *Journal of Research in Personality, 37*, 504-528.

Hall, C. C., Goren, A., Chaiken, S., & Todorov, A. (2009). Shallow cues with deep effects: Trait judgments from faces and voting decisions. In E. Borgida, J. L. Sullivan, & C. M. Federico (Eds.), *The political psychology of democratic citizenship*. Oxford University Press.

Hamermesh, D., & Biddle, J. (1994). Beauty and the labor market. *The American Economic Review, 84*, 1174-1194.

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

Kenny, D. A., Horner, C., Kashy, D. A., & Chu, L. C. (1992). Consensus at zero acquaintance: Replication, behavioral cues, and stability. *Journal of Personality and Social Psychology, 62*, 88-97.

Kim, M. P., & Rosenberg, S. (1980). Comparison of two structural models of implicit personality theory. *Journal of Personality and Social Psychology, 38*, 375-389.

Knutson, B. (1996). Facial expressions of emotion influence interpersonal trait inferences. *Journal of Nonverbal Behavior, 20*, 165-181.

Langlois, J. H., Kalakanis, L., Rubenstein, A. J., Larson, A., Hallam, M., & Smoot, M. (2000). Maxims or myths of beauty? A meta-analytic and theoretical review. *Psychological Bulletin, 126*, 390-423.

LeDoux, J. E. (2000). Emotion circuits in the brain. *Annual Review of Neuroscience, 23*, 155-184.

Locher, P., Unger, R., Sociedade, P., & Wahl, J. (1993). At first glance: Accessibility of the physical attractiveness stereotype. *Sex Roles, 28*, 729-743.

Lundqvist, D., Flykt, A., & Öhman, A. (1998). *The Karolinska directed emotional faces* [Database of standardized facial images]. Psychology Section, Department of

Clinical Neuroscience, Karolinska Hospital, S-171 76 Stockholm, Sweden.

Macrae, C. N., Quinn, K. A., Mason, M. F., & Quadflieg, S. (2005). Understanding others: The face and person construal. *Journal of Personality and Social Psychology, 89*, 686-695.

Mazur, A., Mazur, J., & Keating, C. (1984). Military rank attainment of a West Point class: Effects of cadets' physical features. *American Journal of Sociology, 90*, 125-150.

Montepare, J. M., & Dobish, H. (2003). The contribution of emotion perceptions and their overgeneralizations to trait impressions. *Journal of Nonverbal Behavior, 27*, 237-254.

Montepare, J. M., & Zebrowitz, L. A. (1998). Person perception comes of age: The salience and significance of age in social judgments. *Advances in Experimental Social Psychology, 30*, 93-161.

Mueller, U., & Mazur, A. (1996). Facial dominance of West Point cadets as a predictor of later military rank. *Social Forces, 74*, 823-850.

Olson, I. R., & Marshuetz, C. (2005). Facial attractiveness is appraised in a glance. *Emotion*, *5*, 498-502.

Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences of the USA, 105*, 11087-11092.

Osgood, C. E., Suci, G. I., & Tennenbaum, P. H. (1957). *The measurement of meaning*. Urbana: University of Illinois Press.

Pasley, B. N., Mayes, L. C., & Schultz, R. T. (2004). Subcortical discrimination of unperceived objects during binocular rivalry. *Neuron, 42*, 163-172.

Penton-Voak, I. S., Pound, N., Little, A. C., & Perrett, D. I. (2006). Personality judgments from natural and composite facial images: More evidence for a "kernel of truth" in social perception. *Social Cognition, 24*, 607-640.

Phelps, E. A., & LeDoux, J. E. (2005). Contributions of the amygdala to emotion processing: From animal models to human behavior. *Neuron, 48*, 175-187.

Phillips, M. L., Williams, L. M., Heining, M., Herba, C. M., Russell, T. et al. (2004). Differential neural responses to overt and covert presentations of facial expressions of fear and disgust. *NeuroImage*, *21*, 1484-1496.

Rosenberg, S., Nelson, C., & Vivekananthan, P. S. (1968). A multidimensional approach to the structure of personality impressions. *Journal of Personality and Social Psychology, 9*, 283-294.

Rubenstein, A. J. (2005). Variation in perceived attractiveness: Differences between dynamic and static faces. *Psychological Science, 16*, 759-762.

Said, C. P., Sebe, N., & Todorov, A. (2009). *Structural resemblance to emotional expressions predicts evaluation of emotionally neutral faces.* (Manuscript under review)

Secord, P. F. (1958). Facial features and inference processes in interpersonal perception. In R. Tagiuri & L. Petrullo (Eds.), *Person perception and interpersonal behavior* (pp. 300-315). Stanford, CA: Stanford University Press.

Singular Inversions. (2006). *FaceGen 3.1 Full SDK Documentation*. http://facegen.com

Steiger, J. H. (1980). Tests for comparing elements of a correlation matrix. *Psychological Bulletin, 87*, 245-251.

Todorov, A. (2008). Evaluating faces on trustworthiness: An extension of systems for recognition of emotions signaling approach/avoidance behaviors. In A. Kingstone & M. Miller (Eds.), The year in cognitive neuroscience 2008, *Annals of the New York Academy of Sciences, 1124*, 208-224.

Todorov, A. (In press). Evaluating faces on social dimensions. In A. Todorov, S. T. Fiske, & D. Prentice (Eds), *Social neuroscience: Toward understanding the underpinnings of the social mind*. New York: Oxford University Press.

Todorov, A., Baron, S., & Oosterhof, N. N. (2008). Evaluating face trustworthiness: A model based approach. *Social, Cognitive, & Affective Neuroscience, 3*, 119-127.

Todorov, A., & Engell, A. (2008). The role of the amygdala in implicit evaluation of emotionally neutral faces. *Social, Cognitive, & Affective Neuroscience*.

Todorov, A., Mandisodza, A. N., Goren, A., & Hall, C. C. (2005). Inferences of competence from faces predict election outcomes. *Science, 308*, 1623-1626.

Todorov, A., Said, C. P., Engell, A. D., & Oosterhof, N. N. (2008). Understanding evaluation of faces on social dimensions. *Trends in Cognitive Sciences*.

Todorov, A., & Uleman, J. S. (2003). The efficiency of binding spontaneous trait inferences to actor's faces. *Journal of Experimental Social Psychology, 39*, 549-562.

Uleman, J. S., Blader, S., & Todorov, A. (2005). Implicit impressions. In R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 362-392). New York: Oxford University Press.

Vazire, S., & Gosling, S. D. (2004). E-perceptions: Personality impressions based on personal websites. *Journal of Personality and Social Psychology, 87*, 123-132.

Whalen, P. J. (1998). Fear, vigilance, and ambiguity: Initial neuroimaging studies of the human amygdala. *Current Directions in Psychological Science, 7*, 177-188.

Whalen P. J., Kagan, J., Cook, R. G., Davis, F. C., Kim, H., Polis S., McLaren, D. G., Somerville, L.H., McLean, A.A., Maxwell, J.S. and Johnstone, T. (2004). Human amygdala responsivity to masked fearful eye whites. *Science, 306*, 2061.

Whalen, P. J., Rauch, S. L., Etcoff, N. L., McInerney, S. C., Lee, M. B., & Jenike, M. A. (1998). Masked presentations of emotional facial expression modulate amygdala activity without explicit knowledge. *Journal of Neuroscience, 18*, 411-418.

Wiens, S. (2006). Current concerns in visual masking. *Emotion, 6*, 675-680.

Williams, M. A., Morris, A. P., McGlone, F., Abbott, D. F., & Mattingley, J. B. (2004). Amygdala responses to fearful and happy facial expressions under conditions of binocular rivalry. *The Journal of Neuroscience, 24*, 2898-2904.

Willis, J., & Todorov, A. (2006). First impressions: Making up your mind after 100 ms exposure to a face. *Psychological Science, 17*, 592-598.

Winkielman, P., Berridge, K. C., & Wilbarger, J. L. (2005). Unconscious affective reactions to masked happy versus angry faces influence consumption behavior and judgments of value. *Personality and Social Psychology Bulletin, 31*, 121-135.

Winston, J., Strange, B., O'Doherty, J., & Dolan, R. (2002). Automatic and intentional brain responses during evaluation of trustworthiness of face. *Nature Neuroscience, 5,* 277-283.

Zebrowitz, L. A. (1999). *Reading faces: Window to the soul?* Boulder, CO: Westview Press.

Zebrowitz, L. A., Andreoletti, C., Collins, M. A., Lee, S. Y., & Blumenthal, J. (1998). Bright, bad, babyfaced boys: Appearance stereotypes do not always yield self-fulfilling prophecy effects. *Journal of Personality and Social Psychology, 75*, 1300-1320.

Zebrowitz, L. A., Collins, M. A., & Dutta, R. (1998). The relationship between appearance personality across the life span. *Personality and Social Psychology Bulletin*, *24*, 736-749.

Zebrowitz, L. A., & Montepare, J. M. (2006). The ecological approach to person perception: Evolutionary roots and contemporary offshoots. In M. Schaller, J.A. Simpson, & D.T. Kenrick (Eds.), *Evolution and social psychology*. New York: Psychology Press.

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

Zebrowitz, L. A., Voinescu, L., & Collins, M. A. (1996). "Wide-eyed" and "crooked-faced": Determinants of perceived and real honesty across the life span. *Personality and Social Psychology Bulletin*, *22*, 1258-1269.