

Shared Perceptual Basis of Emotional Expressions and Trustworthiness Impressions From Faces

Nikolaas N. Oosterhof and Alexander Todorov
Princeton University

Using a dynamic stimuli paradigm, in which faces expressed either happiness or anger, the authors tested the hypothesis that perceptions of trustworthiness are related to these expressions. Although the same emotional intensity was added to both trustworthy and untrustworthy faces, trustworthy faces who expressed happiness were perceived as happier than untrustworthy faces, and untrustworthy faces who expressed anger were perceived as angrier than trustworthy faces. The authors also manipulated changes in face trustworthiness simultaneously with the change in expression. Whereas transitions in face trustworthiness in the direction of the expressed emotion (e.g., high-to-low trustworthiness and anger) increased the perceived intensity of the emotion, transitions in the opposite direction decreased this intensity. For example, changes from high to low trustworthiness increased the intensity of perceived anger but decreased the intensity of perceived happiness. These findings support the hypothesis that changes along the trustworthiness dimension correspond to subtle changes resembling expressions signaling whether the person displaying the emotion should be avoided or approached.

Keywords: face perception, emotions, overgeneralization, trustworthiness, dynamic stimuli

Accurate face perception is critical for successful social interaction. People use dynamic changes in the face to understand the mental states of others and invariant facial features to identify other people (Haxby, Hoffman, & Gobbini, 2000). People also form rapid trait impressions such as trustworthiness and aggressiveness from facial appearance (Bar, Neta, & Linz, 2006; Todorov, Pakrashi, & Oosterhof, in press; Willis & Todorov, 2006). Although these impressions are not necessarily accurate, they affect important social outcomes. For example, the presence of Afro-American features predicts criminal-sentencing decisions (Blair, Judd, & Chapleau, 2004), and competent appearance predicts electoral success (Ballew & Todorov, 2007; Todorov, Mandisodza, Goren, & Hall, 2005).

Although people make multiple trait inferences from emotionally neutral faces, these inferences are highly correlated with each other. For example, principal components analyses (PCA) of judgments on trait dimensions that are used to spontaneously characterize faces show that the first two principal components (PCs) account for more than 80% of the variance of the judgments (Oosterhof & Todorov, 2008). The first PC, which accounts for more than 60% of the variance, reflects the valence evaluation of

faces and the second PC reflects their dominance evaluation.¹ In other words, evaluation of faces on social dimensions can be described within a simple two-dimensional space. It is interesting that, in PCAs of two different sets of faces—natural and computer-generated—out of all trait judgments, trustworthiness judgments were the most highly correlated with the first PC ($r > .91$) and practically uncorrelated with the second PC ($|r| < .06$). These findings suggest that these judgments are an excellent approximation of general valence evaluation of faces (Todorov, 2008).

What is the source of valence evaluation of faces? According to the emotion overgeneralization hypothesis (Knutson, 1996; Montepare & Dobish, 2003; Todorov, 2008; Zebrowitz & Montepare, 2008), resemblance of neutral faces to emotional expressions is perceived as indicating the trait attributes or behavioral tendencies associated with these emotions. For example, even when people pose for neutral expressions, their expressions may convey specific emotional states to others (Malatesta, Fiore, & Messina, 1987), and these states can be attributed to personality dispositions. In fact, trustworthiness judgments of emotionally neutral faces are highly negatively correlated with judgments of anger and highly positively correlated with judgments of happiness of the faces (Todorov & Duchaine, 2008; Winston, Strange, O'Doherty, & Dolan, 2002), suggesting that these judgments are based on subtle facial cues that resemble expressions signaling whether the person should be avoided (anger e.g., Adams, Ambady, Macrae, & Kleck, 2006; Marsh, Ambady, & Kleck, 2005) or can be approached (happiness).

Nikolaas N. Oosterhof and Alexander Todorov, Department of Psychology and the Center for the Study of Brain, Mind and Behavior at Princeton University.

We thank Valerie Loehr for her help with the data collection. This research was supported by National Science Foundation Grant BCS-0446846. Nikolaas Oosterhof was supported by a Huygens Scholarship offered by the Netherlands organization for international cooperation in higher education (NUFFIC).

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

¹ All positive trait judgments (e.g., trustworthy, emotionally stable, attractive) had positive loadings on the first PC and all negative judgments (e.g., mean, weird) had negative loadings on this component. The judgments with highest loadings on the second PC were dominant, aggressive, and confident.

Recently, using a statistical model in which faces are represented as points in a multidimensional space, we modeled trustworthiness judgments and obtained evidence consistent with the hypothesis that these judgments are based on cues signaling approach/avoidance behavior (Oosterhof & Todorov, 2008). Specifically, participants were asked to make trustworthiness judgments of emotionally neutral, computer-generated faces. Based on the mean judgments, we built a dimension optimal for changing face trustworthiness in the multidimensional face space (a vector that is a linear combination of the vectors representing face shape).² A subsequent behavioral study confirmed that judgments of faces generated by the trustworthiness dimension corresponded to the trustworthiness predicted by the model.

More important, although the input for the model of this dimension was based on trustworthiness judgments of emotionally neutral faces, exaggerating faces in the negative direction of the trustworthiness dimension increased attributions of anger, whereas exaggerating faces in the positive direction increased attributions of happiness. For example, in a study in which participants categorized faces as neutral or as expressing one of the basic emotions, the dominant categorization, and the only one that was significantly higher than chance, was angry for extremely exaggerated faces in the negative direction ($-8 SD$) and happy for extremely exaggerated faces in the positive direction ($+8 SD$). The most visible changes in the faces were in the eyes and mouth regions. Moving toward the negative end of the dimension, the eyebrows become more \vee -shaped and the mouth more \cap -shaped. In contrast, moving toward the positive end, the eyebrows become more \wedge -shaped and the mouth more \cup -shaped. These simultaneous changes in eyebrows and mouth shape mimic the action units that underlie expressions of anger and happiness (Ekman & Friesen, 1978).

Here, we sought additional evidence for the hypothesis that perceptions of trustworthiness and expressions of anger and happiness are related. Specifically, we tested whether changes in structural features that signal trustworthiness affect the perception of these expressions. We used a dynamic stimuli paradigm (see Figure 1), in which faces expressed either happiness or anger, because this paradigm allowed us to control face trustworthiness within a single trial by gradually morphing trustworthy into untrustworthy faces and vice versa. To the extent that changes along the trustworthiness dimension are related to changes in facial structure resembling expressions of anger and happiness (e.g., low-to-high trustworthiness would correspond to angry-to-happy), changes in trustworthiness should modulate the perception of these emotions.

To test this hypothesis, we first selected trustworthy- and untrustworthy-looking faces based on trustworthiness judgments of a large set of computer-generated, emotionally neutral faces. Then, we created animations in which the neutral faces changed to either smiling or angry faces (Figure 1a). In the baseline condition, the same face expressed one of the emotions. Consistent with the shared signal hypothesis (Adams & Kleck, 2005), which posits that facial cues that are congruent with an expression (e.g., direct gaze and anger) intensify the perception of the expression, we expected that a) untrustworthy faces expressing anger should be perceived as angrier than trustworthy faces expressing the same emotion; and b) trustworthy faces expressing happiness should be perceived as happier than untrustworthy faces expressing the same emotion.

More important, as we noted above, is that we also manipulated changes in the trustworthiness of the faces during the course of the computer animation (Figure 1b). In incongruent animations, the trustworthiness of the face changed simultaneously with the change in

emotional expression. If changes along the trustworthiness dimension are related to changes in expressions, changes in trustworthiness in the direction of the expressed emotion should increase the intensity of the perceived emotion, and transitions in the opposite direction should decrease this intensity. For example, when an untrustworthy face changes into a trustworthy face, the same happy expression should be perceived as happier than when a trustworthy face changes into another trustworthy face. In contrast, when a trustworthy face changes into an untrustworthy face, the same happy expression should be perceived as less happy than when an untrustworthy face changes into another untrustworthy face. To rule out the possibility that such effects could be simply due to additional motion caused by the change of the face identity, we created congruent animations, in which the identity of the face changed but the trustworthiness was kept constant.

Method

Participants

Sixty undergraduate students from Princeton University participated for partial course credit. Twenty-one participated in a study for selection of trustworthy and untrustworthy faces, and 39 participated in the dynamic stimuli study.

Facial stimuli

We used computer-generated faces created by the Facegen Modeler program version 3.1 (Singular Inversions, 2006; <http://www.facegen.com>). This program uses a data-driven statistical model based on 3D laser scans of faces (Banz & Vetter, 1999) and allows for the generation of novel faces. Using procedures described in Oosterhof and Todorov (2008), we randomly generated 96 emotionally neutral Caucasian faces that were not manipulated on trustworthiness. The faces were rated on a scale from 1 (untrustworthy) to 8 (trustworthy), ($n = 21$, Cronbach's $\alpha = .78$). Out of the 96 faces, we selected five low-trustworthy and five high-trustworthy faces to create computer animations. The respective means for the untrustworthy and trustworthy faces were 3.33 ($SD = 0.34$) and 5.12 ($SD = 0.36$). We used only male faces because bald male faces look more natural than bold female faces (see Figure 1).

Animations

Animations consisted of 51 frames of a face that were created using an automated procedure in Facegen. Facegen supports adding several emotional expressions to any face and the expression strengths can be set anywhere between 0% and 100%. We added expressions using Facegen's Anger and SmileOpen (happy) expression controls. Each animation started with a neutral face at frame 1 to which an emotion was added linearly to either 25% or 50% at frame 51. Thus, there were four emotion conditions: "weak happy" (Smile = 0...25%), "medium happy" (Smile = 0...50%), "weak angry" (Anger = 0...25%), and "medium angry" (Anger = 0...50%). We used maximum emotion strength of 50% (rather than 100%) because we

² In subsequent research, we built a valence dimension in the multidimensional face space based on the first PC derived from a PCA of nine trait judgments. This dimension was very similar to the trustworthiness dimension. One SD change on the former corresponded to 0.96 SD change on the latter.

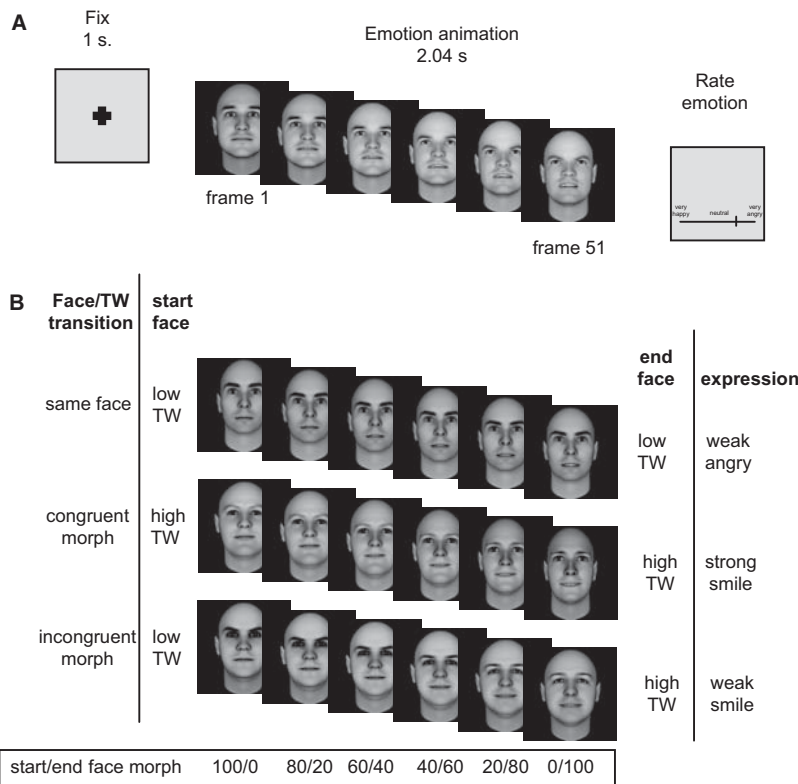


Figure 1. (A) An example of a single experimental trial. During the course of this animation, the expression of the face changes from neutral to angry and the initial trustworthy face morphs into an untrustworthy face. (B) Three other examples of trials representing different experimental conditions.

assumed that modulatory effects of trustworthiness would be stronger for subtler expressions of emotions.

We created three types of animations: same face, congruent trustworthiness morph, and incongruent trustworthiness morph (see Figure 1b for examples). In the same face animations, an emotional expression was added to a single face. We created 40 same face animations (10 faces \times 2 emotions \times 2 levels of emotion strength). The morphed animations were based on two different faces (a start face and an end face) that either matched or mismatched on trustworthiness. During the animation, the target face gradually morphed from the start to the end face (100% start face to 100% end face with 2% change per frame transition), while at the same time an emotional expression was added. Both the start and end faces could be either high or low on trustworthiness, resulting in four trustworthiness conditions for morphed animations. Congruent morphed animations had the same trustworthiness level for the start and end faces (low-to-low and high-to-high), and served as a control for the effects of changes in facial identity on perception of emotions. Incongruent morphed animations had different trustworthiness levels for the start and end faces (high-to-low and low-to-high). Each face was randomly paired with four other faces (two high trustworthy and two low trustworthy) with the constraint that each face was used equally often as a start face and as an end face (i.e., used in 8 different animations: in 4 as a start face and in 4 as an end face). This procedure created 20 congruent face pairs (10 with low and 10 with high trustworthiness: within each group of 10 pairs, to 5 we added weak emotions and to 5

we added strong emotions) and 20 incongruent pairs. Each of these pairs was presented with both happy and angry expressions, resulting in 40 congruent and 40 incongruent animations.

Procedures

The experimental paradigm was implemented in Java 1.5 and run on standard PCs. Participants were informed that they would see animations with facial expressions, in which each face turns from neutral to either happy or angry and that on some trials the face itself may change. They were instructed to judge how happy or angry the end face becomes over the course of the animation.

First, participants did a practice trial to get accustomed with the procedure, and then rated each of the 120 animations. The order of the trials was randomized for each participant. All stimuli were presented on a black square rectangle (400 \times 400 pixels) surrounded by a gray background. On each trial, a white fixation cross was presented for 1 s, followed by a facial expression animation for 2.04 s (25 frames per second).³ Then the stimulus was hidden (gray screen) and the participant rated the emotional expression on a continuous slider scale

³ A potential limitation of our study is that we used relatively long expression durations. This was necessary to allow for smooth transition between faces in the congruent and incongruent animations. However, given that the duration of expressions affects perception of emotions (e.g., Krumbhauer et al., 2007), this is one of the parameters that should be controlled in future studies.

with labels “very happy” (left), “neutral” (center), and “very angry” (right). For the analysis, the scale was transformed into 100 units, ranging from -50 (very angry) to 0 (neutral) to $+50$ (very happy). No time limit was imposed on the participant. By pressing a “next” button the participant continued to the next trial (see Figure 1a).

The overall design was a 2 (Emotion: angry vs. happy) \times 2 (Strength of emotion: weak vs. medium) \times 2 (Face trustworthiness of end face: trustworthy vs. untrustworthy) \times 3 (Morph: same face vs. congruent vs. incongruent) within subject design.

Preliminary analysis

Not surprisingly, a 2 (Emotion: angry vs. happy) \times 2 (Strength of emotion: weak vs. medium) \times 2 (Trustworthiness: trustworthy vs. untrustworthy) \times 3 (Morph: same face vs. congruent vs. incongruent) repeated measures ANOVA revealed a large main effect of emotion, $F(1, 38) = 335.82, p < .001, \eta^2 = .90$. Although this effect was qualified by several interactions, for every simple effect of the emotion factor (e.g., untrustworthy faces with weak happy emotions in the congruent morph condition vs. untrustworthy faces with weak angry emotions in the congruent morph condition), faces with added happy emotions were perceived as significantly happier than faces with added angry emotions, $t_{\min}(38) = 2.62, p < .012, t_{\max}(38) = 22.55$. These findings show that participants reliably identified the emotions of happiness and anger.

In the main analyses, we analyze the data separately for faces with added happy emotions and faces with added angry emotions because of the specific predictions for these conditions and the significant interactions that qualify the main effect. Further, for simplicity, we collapse across the strength of the emotion manipulation. As expected, faces with higher intensity of the emotion were perceived as expressing more strongly the specific emotion, $F(1, 38) = 331.33, p < .001, \eta^2 = .90$, for the interaction of strength of emotion and type of emotion. Separate analyses at the four combinations of type of emotion \times strength of emotion showed that the pattern of responses was the same, and that in all cases the predicted interaction of main theoretical interest—Trustworthiness \times Morph—was significant, $\eta^2_{\min} = .52$.

Results

Perception of happy emotions

As shown in Figure 2, the expressions of trustworthy faces ($M = 15.4, SD = 6.0$) were perceived as happier than the expressions of untrustworthy faces ($M = 3.3, SD = 4.5$), $F(1, 38) = 158.18, p < .001, \eta^2 = .81$. This difference was augmented when the transition in identity involved a change in trustworthiness, as reflected by a significant interaction of trustworthiness and morph, $F(1, 38) = 114.80, p < .001, \eta^2 = .75$. When an untrustworthy face changed into a trustworthy face, the same happy expression was perceived as happier than when a trustworthy face changed into another trustworthy face, $t(38) = 6.55, p < .001$, or when there was no change in the identity of the face, $t(38) = 8.26, p < .001$. Similarly, when a trustworthy face changed into an untrustworthy face, the same happy expression was perceived as less happy than when an untrustworthy face changed into another untrustworthy face, $t(38) = 10.89, p < .001$, or when there was no change in the identity of the face, $t(38) = 12.13, p < .001$.

Perception of angry emotions

As shown in Figure 2, the expressions of untrustworthy faces ($M = -17.7, SD = 6.2$) were perceived as angrier than the expressions of trustworthy faces ($M = -0.4, SD = 4.9$), $F(1, 38) = 329.81, p < .001, \eta^2 = .90$. As in the case of expressions of happiness, this difference was augmented when the transition in identity involved a change in trustworthiness, as reflected by a significant interaction of trustworthiness and morph, $F(1, 38) = 85.73, p < .001, \eta^2 = .69$. When a trustworthy face changed into an untrustworthy face, the same angry expression was perceived as angrier than when an untrustworthy face changed into another untrustworthy face, $t(38) = 5.62, p < .001$, or when there was no change in the identity of the face, $t(38) = 8.22, p < .001$. Similarly, when an untrustworthy face changed into a trustworthy face, the same angry expression was perceived as less angry than when a trustworthy face changed into another trustworthy face, $t(38) = 9.23, p < .001$, or when there was no change in the identity of the face, $t(38) = 8.62, p < .001$.

Discussion

Using a dynamic stimuli paradigm, we showed that face trustworthiness modulates the intensity of perceived emotions. Although the same emotional intensity was added to both trustworthy and untrustworthy faces, we found that a) trustworthy faces who expressed happiness were perceived as happier than untrustworthy faces who expressed the same emotion; and b) untrustworthy faces who expressed anger were perceived as angrier than trustworthy faces who expressed the same emotion.

Moreover, whereas transitions in face trustworthiness in the direction of the expressed emotion increased the perceived intensity of the final emotion, transitions in the opposite direction decreased the perceived intensity of the emotion (see Figure 2). For example, changes from high to low trustworthiness increased the intensity of perceived anger but decreased the intensity of perceived happiness. Similarly, changes from low to high trustworthiness increased the intensity of perceived happiness but decreased the intensity of perceived anger. These findings provide additional evidence that changes along the trustworthiness dimension correspond to subtle changes in features resembling expressions signaling whether the person displaying the emotion should be approached or avoided (Todorov, 2008).

We focused on expressions of anger and happiness because our prior modeling work showed that changes along the trustworthiness dimension primarily affected perceptions of anger and happiness but not perceptions of other basic emotions (Oosterhof & Todorov, 2008). However, these findings do not rule out the possibility that resemblance of neutral faces to other expressions such as fear can affect trustworthiness judgments. One hypothesis is that the extent to which other expressions affect trustworthiness judgments is a function of their similarity to expressions of anger and happiness.

We should also note that resemblance of neutral faces to expressions of anger and happiness is not the only source of inferences of trustworthiness. Other sources include femininity/masculinity (Oosterhof & Todorov, 2008), facial maturity (Montepare & Zebrowitz, 1998), physical similarity to the self (DeBruine, 2005), and possibly facial texture. For example, both




Figure 2. Valence ratings of emotional expressions as a function of the type of emotion, the trustworthiness of the face, and the morphing condition: same face with no change in identity; congruent morph with no change in face trustworthiness but change in identity; and incongruent morph with changes in both face trustworthiness and identity. The ratings were made on a continuous slider ranging from angry to neutral to happy. The error bars show standard errors of the means.

feminine and babyfaced faces are judged as more trustworthy than masculine and mature faces, respectively.

Our research complements previous studies that have found that changes in expression of emotions affect trait impressions (Knutson, 1996; Montepare & Dobish, 2003). Here, using morphing, we showed that changes in structural features of faces could affect the perception of emotions. To the extent that there is a correspondence between facial cues associated with trait judgments and facial cues signaling emotions, it should be possible to observe mutual bidirectional effects on perception. Expressions can affect trait judgments, and facial appearance can affect perception of emotions.

To summarize, based on our prior studies, we argued that valence is the main dimension along which emotionally neutral faces are evaluated and that trustworthiness judgments are an excellent approximation of valence evaluation (Oosterhof & Todorov, 2008; Todorov, 2008). The valence evaluation of stimuli triggers automatic approach/avoidance responses (Chen & Bargh, 1999), and there is evidence that angry expressions trigger automatic avoidance responses (Adams et al., 2006; Marsh et al., 2005). Consistent with this perspective, the current findings and our prior computer modeling findings (Oosterhof & Todorov, 2008), suggest a shared basis of perceptions of face trustworthiness and expressions of anger and happiness. The modeling findings showed that variations in face trustworthiness are based on subtle resemblance of neutral faces to expressions of anger and happiness. The current findings showed that variations in face trustworthiness affect the perception of these emotions.

The shared perceptual basis hypothesis can account for dissociations between processing of facial identity and evaluation of face trustworthiness. Todorov and Duchaine (2008) showed that developmental prosopagnosics who were severely impaired in both perception of and memory for facial identity were nevertheless able to make normal trustworthiness judgments from novel faces. Given that there are dissociations between recognition of facial identity and recognition of emotional expressions in devel-

opmental prosopagnosics (Duchaine, Parker, & Nakayama, 2003), it is possible that the prosopagnosics' preserved processing of face trustworthiness is based on preserved processing of emotional expressions. That is, prosopagnosics with normal recognition of expressions of anger and happiness should perform normally with trustworthiness judgments whereas those with expression deficits should also have trustworthiness judgments deficits. Finally, the shared perceptual basis hypothesis also suggests that trustworthiness judgments may be used as a subtle test of preserved emotional processing of faces. These hypotheses remain to be tested.

References

- Adams, R. B., Ambady, N., Macrae, N., & Kleck, R. E. (2006). Emotional expressions forecast approach-avoidance behavior. *Motivation & Emotion, 30*, 179–188.
- Adams, R. B., & Kleck, R. E. (2005). Effects of direct and averted gaze on the perception of facially communicated emotion. *Emotion, 5*, 3–11.
- Ballew, C. C., & Todorov, A. (2007). Predicting political elections from rapid and unreflective face judgments. *Proceedings of the National Academy of Sciences, USA, 104*, 17948–17953.
- Bar, M., Neta, M., & Linz, H. (2006). Very first impressions. *Emotion, 6*, 269–278.
- 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. *Proceedings of the 26th annual conference on computer graphics and interactive techniques*, (pp. 187–194). Los Angeles: Addison Wesley Longman.
- Chen, M., & Bargh, J. A. (1999). Consequences of automatic evaluation: Immediate behavioral predispositions to approach or avoid the stimulus. *Personality and Social Psychology Bulletin, 25*, 215–224.
- DeBruine, L. M. (2005). Trustworthy but not lust-worthy: Context specific effects of facial resemblance. *Proceedings of the Royal Society, Series B, 272*, 919–922.
- Duchaine, B. C., Parker, H., & Nakayama, K. (2003). Normal emotion recognition in a developmental prosopagnosic. *Perception, 32*, 827–838.

- Ekman, P., & Friesen, W. (1978). *The facial action coding system: A technique for the measurement of facial movement*. San Francisco: Consulting Psychologists Press, Inc.
- Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences*, 4, 223–233.
- Knutson, B. (1996). Facial expressions of emotion influence interpersonal trait inferences. *Journal of Nonverbal Behavior*, 20, 165–181.
- Krumhuber, E., Manstead, A. S. R., Kappas, A., Cosker, D., Marshall, D., & Rosin, P. L. (2007). Facial dynamics as indicators of trustworthiness and cooperative behavior. *Emotion*, 7, 730–735.
- Malatesta, C. Z., Fiore, M. J., & Messina, J. J. (1987). Affect, personality, and facial expressive characteristics of older people. *Psychology and Aging*, 2, 64–69.
- Marsh, A. A., Ambady, N., & Kleck, R. E. (2005). The effects of fear and anger facial expressions on approach- and avoidance-related behaviors. *Emotion*, 5, 119–124.
- 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.
- Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences, USA*, 105, 11087–11092.
- Singular Inversions (2006). FaceGen 3.1 Full Software Development Kit Documentation. <http://www.facegen.com>. Retrieved on June 5th, 2007.
- 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., & Duchaine, B. (2008). Reading trustworthiness in faces without recognizing faces. *Cognitive Neuropsychology*, 25, 395–410.
- 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., Pakrashi, M., & Oosterhof, N. N. (In press). Evaluating faces on trustworthiness after minimal time exposure. *Social Cognition*.
- Willis, J., & Todorov, A. (2006). First impressions: Making up your mind after 100 ms exposure to a face. *Psychological Science*, 17, 592–598.
- Winston, J. S., Strange, B. A., O'Doherty, J., & Dolan, R. J. (2002). Automatic and intentional brain responses during evaluation of trustworthiness of faces. *Nature Neuroscience*, 5, 277–283.
- Zebrowitz, L. A., & Montepare, J. M. (2008). Social psychological face perception: Why appearance matters. *Social and Personality Psychology Compass*, 2, 1497–1517.

Received March 18, 2008

Revision received October 13, 2008

Accepted October 14, 2008 ■