



Adaptation to facial trustworthiness is different in female and male observers[☆]



Joanna Wincenciak^a, Milena Dzhelyova^b, David I. Perrett^b, Nick E. Barraclough^{c,*}

^aDepartment of Psychology, University of Hull, Cottingham Road, Hull HU6 7RX, UK

^bSchool of Psychology, University of St Andrews, St Mary's Quad, South Street, St Andrews, Fife KY16 9JP, UK

^cDepartment of Psychology, University of York, Heslington, York YO10 5DD, UK

ARTICLE INFO

Article history:

Received 30 October 2012

Received in revised form 14 May 2013

Available online 28 May 2013

Keywords:

Face

Adaptation

Trustworthiness

Gender difference

ABSTRACT

Face adaptation paradigms have been used extensively to investigate the mechanisms underlying the processing of several different facial characteristics including face shape, identity, view and emotional expression. Judgements of facial trustworthiness can also be influenced by visual adaptation; to date these (un)trustworthy face aftereffects have only been shown following adaptation to emotional expression and facial masculinity/femininity. In this study we assessed how exogenous test faces and untrustworthy faces influenced the perception of the trustworthiness of subsequent test faces. In a mixed factorial design experiment, we tested the influence of adaptation to female and male faces on the perception of subsequent female and male faces in both female and male observers. In female observers, we found that following adaptation to trustworthy and untrustworthy faces subsequent test faces appeared less like the adapting stimuli. Sex of the adapting and test faces did not have significant influence on these (un)trustworthy face aftereffects. In male observers, however, we found no significant influence of the effect of adaptation on the subsequent perception of face trustworthiness. The clear difference in the visual aftereffects induced in female and male observers indicates the operation of different mechanisms underlying the perception of facial trustworthiness, and future studies should investigate these mechanisms separately in female and male observers.

© 2013 The Authors. Published by Elsevier Ltd. All rights reserved.

1. Introduction

Faces provide detailed visual information about other people's identity, gender, direction of attention, emotion and internal mental states. People also derive complex trait judgements, such as competence (e.g. Ballew & Todorov, 2007) and trustworthiness (Winston et al., 2002) from facial appearance, and such judgements are made rapidly and accurately and can help guide decision making during social interactions. Trustworthiness conveyed by a person's face, for example, is used by observers when making trust decisions and whether to invest money with that person (e.g. van't Wout & Sanfey, 2008). Judgements of facial trustworthiness are highly consistent even with very brief exposure times (Todorov, Pakrashi, & Oosterhof, 2009; Willis & Todorov, 2006), allowing for rapid assessment of useful social information within a single glance.

[☆] This is an open-access article distributed under the terms of the Creative Commons Attribution-NonCommercial-No Derivative Works License, which permits non-commercial use, distribution, and reproduction in any medium, provided the original author and source are credited.

* Corresponding author. Fax: +44 (0)1904 4331.

E-mail address: nick.barraclough@york.ac.uk (N.E. Barraclough).

Face perception, however, is subject to adaptation, where prior exposure to a visual stimulus generates subsequent aftereffects resulting in biases in perception. Face adaptation experiments typically show that adaptation to faces conveying specific facial characteristics result in aftereffects where subsequently presented faces appear less like the adapting face (see Webster & MacLeod, 2011 for a review). For example, adaptation to a male face shifts perception of subsequent gender neutral faces so that they are more often categorised as female faces. Similarly, adaptation to female faces causes equivalent aftereffects, where gender neutral faces are more often categorized as male faces (Webster et al., 2004). Such 'repulsive' aftereffects have been observed following adaptation to many other facial characteristics including face shape (Rhodes et al., 2004; Webster & MacLin, 1999), identity (Leopold et al., 2001), emotional expression (Fox & Barton, 2007), and view (Chen et al., 2010). Thus, many of the judgements we make about other people's faces appear highly dynamic, and dependent upon the immediate prior visual context.

Perception of facial trustworthiness is susceptible to visual adaptation, but so far this has only been demonstrated following adaptation to other facial characteristics. For example, adaptation to angry or happy faces biases evaluation of subsequent neutral faces so that they appear more or less trustworthy respectively;

in contrast, adaptation to fearful faces has no effect on perceived facial trustworthiness (Engell, Todorov, & Haxby, 2010). This research supports proposals that evaluations of facial trustworthiness and judgements of specific emotional expressions derived from faces rely on some common facial features and thus may in part be processed by overlapping neural mechanisms (e.g. Oosterhof & Todorov, 2009; Todorov, Baron, & Oosterhof, 2008; Todorov & Engell, 2008; Todorov et al., 2008).

Furthermore, perception of facial trustworthiness has also been shown to be biased following adaptation to masculine and feminine faces (Buckingham et al., 2006). In this study participants adapted to both masculinised and feminised versions of male faces, and were required to rate the attractiveness or trustworthiness of subsequent test male faces. Adaptation to both masculinised and feminised faces appeared to affect the perceived trustworthiness of subsequent faces. Buckingham et al. argued that this adaptation mechanism may contribute to the development of longer term face preference judgements.

It is not yet clear, however, if adaptation occurs following exposure to faces that differ in perceived levels of trustworthiness. If facial trustworthiness adaptation follows the typical pattern as seen for previously tested face aftereffects where subsequent stimuli appear less like the adapting stimulus, then we would expect to find that adaptation to untrustworthy faces causes subsequent faces to appear more trustworthy, and adaptation to trustworthy faces causes subsequent faces to appear more untrustworthy. Aftereffects resulting from trustworthiness adaptation might have implications for the development of longer term judgments of the trustworthiness of other individuals from daily experiences of faces (e.g. see Carbon & Ditye, 2011).

To investigate whether we adapt to facial trustworthiness, here we compared the effect of adaptation to untrustworthy and trustworthy faces in order to assess whether facial trustworthiness adaptation biased subsequent perception of facial trustworthiness. Previous research has indicated that the perception of facial trustworthiness is contingent of the sex of the face (e.g. Dzhelyova, Perrett, & Jentsch, 2012). We therefore measured the effect of adaptation to both trustworthy and untrustworthy female and male faces and its effect upon the perception of female and male test faces in both female and male observers, all within one factorial design experiment.

2. Methods

2.1. Participants

Participants were University of Hull students and staff (24 females, mean age = 22.5 years, SD = 4.71; 24 males, mean age = 23.9, SD = 6.01), all received course credit or payment for participating. All participants had normal or corrected to normal vision and were naive to the purpose of the study. Experiments were approved by the ethics committee of the Department of Psychology, University of Hull, and performed in accordance with the ethical standards laid down in the 1990 Declaration of Helsinki.

2.2. Stimuli

Face stimuli were obtained from The Perception Lab, University of St Andrews. The original set of stimuli included ninety-nine faces (49 male) of Caucasian students, age range 17–25, displayed on a white background with a neutral expression, minimal makeup and no jewellery, and were horizontally aligned and scaled to the same interpupillary distance. Each face was rated for trustworthiness using a 7-point Likert scale by independent observers. Adapting faces (illustrated in Fig. 1) were prototypes generated by




Fig. 1. Female and male adapting stimuli.

averaging (Rowland & Perrett, 1995) separately the 8 most untrustworthy and the 8 most trustworthy faces of each sex from the bank of 99 images. Test faces were a subset of 50 images (25 females, 25 males, age range: 18–24) from the original set of 99. The images were selected to cover the whole range of trustworthiness ratings.

In order to ensure that adaptation to low-level visual properties did not contribute to any measure of trustworthiness aftereffects, all adapting stimuli (subtended 356×455 pixels) were presented twice the size of the test stimuli (subtended 178×240 pixels). Face aftereffects typically transfer across substantial changes in stimulus size (Leopold et al., 2001; Zhao & Chubb, 2001) and indicate that the source of the face aftereffect is at a high-level object-based representation of the face rather than a lower-level retinotopic image-based representation of the face.

2.3. Design and experimental procedure

The experiment was controlled by a PC running MATLAB 2006a with the Cogent toolbox. Stimuli were displayed in the centre of a 22" screen CRT monitor (Philips 202P40, 1600×1200 pixels, 100 Hz refresh rate). Participant responses were recorded on the number pad of the keyboard.

The experiment consisted of four blocks of 50 trials which followed a standard adaptation procedure. Each block began with a single 30 s long pre-adaptation phase, followed by a testing phase consisting of alternating presentation of 'top-up' adaptation and test stimuli. Between the pre-adaptation and testing phase instructions were briefly (1000 ms) displayed indicating that the testing phase was about to begin. All stimuli were presented on a mid-grey screen throughout.

In each block one of the four prototype faces (male or female \times trustworthy, untrustworthy, see Fig. 1) was used as the adapting stimulus for both the pre-adaptation phase and the testing phase. All 50 test stimuli (25 female, 25 male) were seen once in each of the 4 blocks. Each of the 50 trials during the testing phase consisted of first a 5000 ms presentation of the adapting stimulus, a brief inter-stimulus interval (ISI, duration 500 ms) then the presentation of one of the smaller test stimuli (duration 250 ms). Following the presentation of the test stimulus participants indicated their judgement of the untrustworthiness of the

Table 1

Mean ratings of prototypes and test faces on a 8-point Likert scale, where 1 = very trustworthy, 8 = very untrustworthy.

	Untrustworthy prototype		Trustworthy prototype		Average of 50 test faces		Average of 17 test faces	
	Female faces	Males faces	Female faces	Males faces	Female faces	Males faces	Female faces	Males faces
Female observers	4.6	5.2	2.3	3.5	4.6	5.2	3.9	4.3
Male observers	4.7	5.2	2.4	3.9	4.5	5.1	3.7	4.5

test face using an 8 point Likert scale where 1 = very trustworthy and 8 = very untrustworthy. Once a response was registered the screen remained grey for a further 1000 ms before the beginning on the next trial. Test stimuli were presented in a random order within each block. After completing each block of 50 trials, participants were offered a short break before proceeding to a subsequent block. Order of testing blocks was counterbalanced using a Latin square method in order to eliminate position effects for the different adapting stimuli.

In order to ensure that the faces used as test stimuli lay on the continuum of (un)trustworthiness between the trustworthy prototypes and the untrustworthy prototypes we performed an additional experiment to rate all of the used in the adaptation experiment. A second set of independent observers (14 females, mean age = 20.5 years, SD = 3.40; 14 males, mean age = 23.9, SD = 5.46) rated both the female and male untrustworthy and trustworthy prototypes together with the 50 individual test faces on a 1–8 Likert scale (1 = very trustworthy, 8 = very untrustworthy). Each original face was shown once and prototype faces were shown 5 times; faces were shown in a random order.

3. Results

Mean trustworthiness ratings for the 4 prototype faces and the 50 test faces were calculated for each participant and then averaged across participants (see Table 1). Surprisingly many of the original test faces were seen as more untrustworthy than the untrustworthy prototypes, this may have resulted from the averaging process used in the generation of our untrustworthy prototypes. Average faces, for example, can appear more attractive than the originals from which they are made (DeBruine et al., 2007), and perception of trustworthiness is highly correlated with attractiveness (Oosterhof & Todorov, 2008).

We wanted to analyse the ratings of those test faces where adaptation to the trustworthy and untrustworthy prototypes would generate opposite aftereffects. We therefore assessed the effect of adaptation on those test faces who's rating of trustworthiness was between the values for the trustworthy and untrustworthy prototypes, when rated by both female and male observers. In total there were 14 (7 female) test faces that matched this criterion, their mean ratings are also described in Table 1¹.

Ratings of the 14 test stimuli were entered into a 4-way mixed ANOVA with 3 within group factors [adaptation trustworthiness (trustworthy, untrustworthy) × adaptation stimulus sex (female, male) × test stimulus sex (female, male)] and 1 between group factor [participant gender (female, male)]. ANOVA showed that there was a significant effect of adaptation trustworthiness (main effect of adaptation trustworthiness: $F(1,46) = 6.87, p < 0.05, \eta_p^2 = 0.13$),

¹ The additional 36 (18 female) faces that lie outside the continuum between the two types of adapting prototype stimuli were excluded from the main analyses. As these faces were rated as more untrustworthy than even the untrustworthy prototype any interpretation of the aftereffect induced in these stimuli becomes problematic (see discussion section). We therefore only analyse aftereffects in the restricted set of 14 stimuli where data can be interpreted, although the inclusion of all 50 test faces produces a similar finding. We note, however, that face aftereffects have previously been assessed with smaller number of test stimuli than in this study (e.g. Buckingham et al., 2006; DeBruine et al., 2007).

and importantly a significant interaction between adaptation trustworthiness and participant gender ($F(1,46) = 8.88, p < 0.005, \eta_p^2 = 0.16$). We also found a significant main effect of test stimulus sex ($F(1,46) = 31.45, p < 0.001, \eta_p^2 = 0.41$) indicating that male faces were on average judged as more untrustworthy than female faces. No other main effects or interactions were significant (all $F_s < 2.19$, all $p_s > 0.145$, all $\eta_p^2 < 0.05$). As this 4-way ANOVA indicated that aftereffects were significantly different in female and male observers, we analysed aftereffects from these 2 groups of participants separately using two 3-way ANOVAs in order to better interpret and understand the effect of adaptation.

3.1. Female participants

Ratings of test stimuli for female observers were analysed using a 3-way ANOVA [adaptation trustworthiness (trustworthy, untrustworthy) × adaptation stimulus sex (female, male) × test stimulus sex (female, male)]. This analysis showed a significant main effect of adaptation trustworthiness ($F(1,23) = 14.75, p < 0.001, \eta_p^2 = 0.39$), and a main effect of test stimulus sex ($F(1,23) = 13.33, p < 0.001, \eta_p^2 = 0.37$) indicating that female test faces ($M = 3.86, SD = 1.08$) were on average rated as more trustworthy than male test faces ($M = 4.74, SD = 0.87$). No other main effects or interactions were found (all $F < 1.11$, all $p > 0.30$, all $\eta_p^2 < 0.05$).

In order to illustrate the effect of adaptation we calculated aftereffects (ratings following trustworthy adaptation – ratings following untrustworthy adaptation) for each condition (see Fig. 2). All aftereffects were positive on this scale indicating that following adaptation, test faces looked less like the preceding adapting stimuli; this is characteristic of other face aftereffects (see Webster & MacLeod, 2011 for a review).

3.2. Male participants

We analysed ratings of test stimuli measured in male participants in the same way as for female observers by using a separate 3-way ANOVA [adaptation trustworthiness (trustworthy, untrustworthy) × adaptation stimulus sex (female, male) × test stimulus sex (female, male)]. For male participants adaptation trustworthiness had no effect on the rating of test stimuli (main effect of adaptation trustworthiness: $F(1,23) = 0.07, p = 0.79, \eta_p^2 = 0.003$). There was, however, a main effect of test stimulus sex ($F(1,23) = 18.39, p < 0.0001, \eta_p^2 = 0.44$) where female test faces ($M = 3.82, SD = 0.87$) were on average rated as more trustworthy than male test faces ($M = 4.82, SD = 0.91$) across all experimental conditions. No other main effects or interactions were found (all $F < 0.78$, all $p > 0.38$, all $\eta_p^2 < 0.03$). As for female observers we calculated aftereffects (ratings following trustworthy adaptation – ratings following untrustworthy adaptation) for each condition, and these are illustrated in Fig. 3.

4. Discussion

The results of this experiment suggest that adaptation to face trustworthiness generates different aftereffects depending upon the gender of the observer. Female observers showed typically

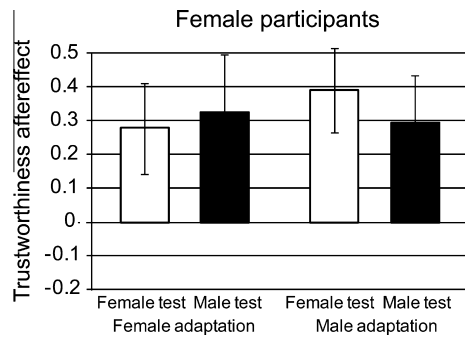


Fig. 2. The effect of adaptation to (un)trustworthiness on faces in female observers. Positive values indicate 'repulsive' aftereffects where test stimuli are judged as less like the adapting stimuli, negative values indicate priming like effects where stimuli are judged as more like the adapting stimuli. Error bars indicate SEM.

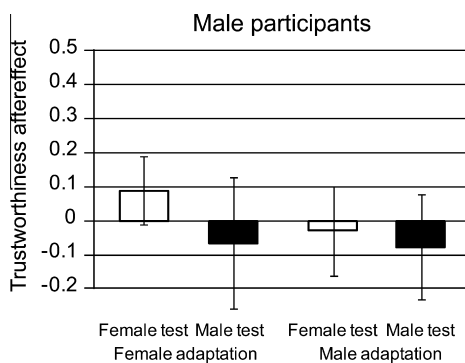


Fig. 3. The effect of adaptation to (un)trustworthiness on faces in male observers. Positive values indicate 'repulsive' aftereffects where test stimuli are judged as less like the adapting stimuli, negative values indicate priming like effects where stimuli are judged as more like the adapting stimuli. Error bars indicate SEM.

characteristic adaptation aftereffects. Here, adaptation to untrustworthy and trustworthy faces resulted in 'repulsive' aftereffects, where test stimuli are judged as less like the adapting stimuli. Similar repulsive aftereffects have been observed many times previously following adaptation to other facial characteristics including shape (Webster & MacLin, 1999), identity (Leopold et al., 2001), emotional expression (Fox & Barton, 2007) and view (Chen et al., 2010). Male observers, in contrast, showed no face adaptation aftereffects. For both female and male observers, our male test stimuli were always judged as more untrustworthy than our female test stimuli (see Table 1). Critically, however, for both female and male observers, there was no interaction between the effect of test stimulus sex and adaptation trustworthiness indicating that the different ratings of these two groups of test stimuli did not influence aftereffects.

Repulsive face aftereffects are thought to involve changes in the sensitivity of cells late in the visual processing hierarchy that selectively code faces (e.g. Chen et al., 2010; Webster & MacLeod, 2011). Prolonged exposure to a face, or faces, with a defining set of characteristics could, in principle, generate a temporary reduction in sensitivity in cells coding those specific facial characteristics (along with other changes in cell tuning functions). Other neural mechanisms coding different facial properties, that don't respond to the adapting face, would retain their sensitivity. One explanation for the facial trustworthiness aftereffects observed here in female observers is that adaptation to untrustworthy faces results in a temporary reduction in the sensitivity of the population of cells that represents untrustworthy faces. The cell population that represents trustworthy faces would have its sensitivity spared. The

cell population response to subsequent test faces would then be biased towards the more trustworthy end of the spectrum, leading to the observed aftereffects where the test faces appear more trustworthy. A similar, but opposite process, could occur following adaptation to trustworthy faces. Face trustworthiness is determined by a number of different facial characteristics (Oosterhof & Todorov, 2008, 2009; Stirrat & Perrett, 2010), and so (un)trustworthiness adaptation is likely to affect multiple populations of neurons coding these various features.

For female observers there was no interaction between adaptation trustworthiness and adaptation stimulus sex or test stimulus sex. This indicates that the aftereffects are similarly sized irrespective of the sex of the adapting and test stimuli (see also Fig. 2). It appears, therefore, that the (un)trustworthy aftereffect in female observers is a general effect and does not take into account the sex of the faces observed, indicating sex-independent coding of facial trustworthiness in female observers.

In contrast to the aftereffects observed in female participants, in male participants adaptation did not produce any repulsive or attractive aftereffects. There was an indication that for male observers that following adaptation, male test faces looked slightly more like the adapting stimuli, an attractive effect similar perhaps to visual object priming. However, there was no significant interaction between adaptation trustworthiness and adaptation stimulus sex or test stimulus sex. A possible explanation for our findings is that the aftereffects that result from adaptation to the trustworthiness of faces by male observers are dominated by other cognitive effects rather than putative (un)trustworthy aftereffects. There is evidence that during judgements of trustworthiness male observers can be less influenced by visual information from the face than female observers. In an investigation of the role of facial width-to-height ratio on perception of trustworthiness, male observers were less influenced by this perceptual information than female observers (Stirrat & Perrett, 2010). In particular subordinate females were most likely to be influenced by facial width-to-height ratio when making judgements of trustworthiness. Stirrat and Perrett's explanation for these findings was that male observers, and especially more dominant males, may be able to effectively ignore facial width-to-height ratios as a cue to another individual's trustworthiness as dominant males often can operate with impunity. For male participants, perhaps small (un)trustworthiness aftereffects are also effectively ignored.

Our finding that female observers show typical adaptation aftereffects to untrustworthy and trustworthy faces whilst male observers show no aftereffects may reflect a different balance in the influence of different processing mechanisms underlying perception of trustworthiness in female and male participants. Although the role of gender of participant in facial trustworthiness perception is most often not studied, there is some electrophysiological evidence that there may be differences in the way facial trustworthiness is processed in female and male observers. In an event related potential (ERP) study of the perception of untrustworthy and trustworthy faces, female observers were faster and more accurate in the perception of facial trustworthiness than male observers (Dzhelyova, Perrett, & Jentsch, 2012). Furthermore, the amplitude of the P1 signal (a positive signal 100 ms after stimulus onset) over the right hemisphere was larger in female than male observers during observation of faces. This sex difference in the P1 signal, however, may relate to face processing *per se* rather than the processing of trustworthiness from faces.

The test stimuli that were used to assess the effect of adaptation to trustworthy and untrustworthy faces were a restricted set of 14 rather than all 50 test stimuli originally shown to participants. This was necessary as the rating of the test stimuli by the set of independent observers, indicated that, against original predictions, many of the test stimuli were rated as even more untrustworthy

than our female and male untrustworthy prototypes. This may have resulted from the averaging process used in our generation of the untrustworthy prototypes. We selected the 8 most untrustworthy faces from the original database of 99 faces to form each of the female and male untrustworthy prototypes. However, the morphing technique to make the prototypes generates average faces (Rowland & Perrett, 1995), and average faces can appear more attractive than the originals from which they are made (discussed in DeBruine et al., 2007). For example, our average untrustworthy prototypes are likely to be more symmetrical, and have smoother skin than many of the test faces. As the perception of trustworthiness is highly correlated with attractiveness (Oosterhof & Todorov, 2008), our untrustworthy female and male prototype faces may have appeared more attractive and thus more trustworthy than many of the potentially more irregular faces in our original set of 50. Because our measure of the effect of adaptation involves the comparison of the effect of adaptation to prototypes on opposite ends of the trustworthiness–untrustworthiness continuum, this makes it problematic to try and analyse the effect of adaptation on the particularly untrustworthy test faces. We have not analysed the differential effect of adaptation to our trustworthy and untrustworthy prototypes that are both perceived as more trustworthy than these particularly untrustworthy test stimuli, as we have no specific predictions about how such adaptation may influence these stimuli.

In conclusion, this study found that aftereffects result from prolonged exposure to untrustworthy and trustworthy faces for female participants but not for male participants, complementing earlier reports of sexual dimorphism in facial trustworthiness. These aftereffects were measured while participants assessed photographs of faces, indicating that in female observers immediate prior visual experience to facial trustworthiness can potentially bias judgments made about the trustworthiness of faces in real-life situations. The clear difference in the visual aftereffects induced in female and male observers indicates the operation of different mechanisms underlying the perception of facial trustworthiness, and future studies should investigate these mechanisms separately in female and male observers.

Acknowledgments

This work was supported by a Grant from the ESRC (RES-062-23-2797); Joanna Wincenciak was supported by an 80th Anniversary Scholarship from the University of Hull.

References

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

- Buckingham, G., DeBruine, L. M., Little, A. C., Welling, L. L. M., Conway, C. A., Tiddeman, B. P., et al. (2006). Visual adaptation to masculine and feminine faces influences generalized preferences and perceptions of trustworthiness. *Evolution and Human Behaviour*, 27, 381–389.
- Carbon, C.-C., & Ditye, T. (2011). Sustained effects of adaptation on the perception of familiar faces. *Journal of Experimental Psychology: Human Perception and Performance*, 37(3), 615–625.
- Chen, J., Yang, H., Wang, A., & Fang, F. (2010). Perceptual consequences of face viewpoint adaptation: Face viewpoint aftereffect, changes of differential sensitivity to face view, and their relationship. *Journal of Vision*, 10(3), 1–11.
- DeBruine, L. M., Jones, B. C., Unger, L., Little, A. C., & Feinberg, D. R. (2007). Dissociating averageness and attractiveness: Attractive faces are not always average. *Journal of Experimental Psychology: Human Perception and Performance*, 33, 1420–1430.
- Dzhelyova, M. P., Perrett, D. I., & Jentzsch, I. (2012). Temporal dynamics of trustworthiness perception. *Brain Research*, 1435, 81–90.
- Engell, A. D., Todorov, A., & Haxby, J. V. (2010). Common neural mechanisms for the evaluation of facial trustworthiness and emotional expressions as revealed by behavioural adaptation. *Perception*, 39, 931–941.
- Fox, C. J., & Barton, J. J. S. (2007). What is adapted in face adaptation? The neural representations of expression in the human visual system. *Brain Research*, 1127, 80–89.
- Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience*, 4(1), 89–94.
- Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences*, 105(32), 11087–11092.
- Oosterhof, N. N., & Todorov, A. (2009). Shared perceptual basis of emotional expressions and trustworthiness impressions from faces. *Emotion*, 9(1), 128–133.
- Rhodes, G., Jeffery, L., Watson, T. L., Jaquet, E., Winkler, C., & Clifford, C. W. G. (2004). Orientation-contingent face aftereffects and implications for face-coding mechanisms. *Current Biology*, 14, 2119–2123.
- Rowland, D. A., & Perrett, D. I. (1995). Manipulating facial appearance through shape and color. *IEEE Computer Graphics and Applications*, 15(5), 70–76.
- Stirrat, M., & Perrett, D. I. (2010). Valid facial cues to cooperation and trust: Male facial width and trustworthiness. *Psychological Science*, 21(3), 349–354.
- Todorov, A., Baron, S. G., & Oosterhof, N. N. (2008). Evaluating face trustworthiness: A model based approach. *Social Cognitive and Affective Neuroscience*, 3, 119–127.
- Todorov, A., & Engell, A. D. (2008). The role of the amygdala in implicit evaluation of emotionally neutral faces. *Social Cognitive and Affective Neuroscience*, 3, 303–312.
- Todorov, A., Pakrashi, M., & Oosterhof, N. (2009). Evaluating faces on trustworthiness after minimal time exposure. *Social Cognition*, 27(6), 813–833.
- Todorov, A., Said, C. P., Engell, A. D., & Oosterhof, N. (2008). Understanding evaluation of faces on social dimensions. *Trends in Cognitive Science*, 12(12), 455–460.
- van't Wout, M., & Sanfey, A. G. (2008). Friend or foe: The effect of implicit trustworthiness judgments in social decision-making. *Cognition*, 108, 796–803.
- Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature*, 428, 557–561.
- Webster, M. A., & MacLeod, D. I. A. (2011). Visual adaptation and face perception. *Philosophical Transactions of the Royal Society*, 366, 1702–1725.
- Webster, M. A., & MacLin, O. H. (1999). Figural aftereffects in the perception of faces. *Psychonomic Bulletin and Review*, 6, 647–653.
- Willis, J., & Todorov, A. (2006). First impressions: Making up your mind after a 100-ms exposure to a face. *Psychological Science*, 17(7), 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(3), 277–283.
- Zhao, L., & Chubb, C. (2001). The size-tuning of the face-distortion after-effect. *Vision Research*, 41, 2979–2994.