

Amygdala Responsivity to High-Level Social Information from Unseen Faces

Jonathan B. Freeman,^{1,2} Ryan M. Stoler,^{1,2} Zachary A. Ingbretsen,^{2,3} and Eric A. Hehman^{1,2}

¹Department of Psychology, New York University, New York, New York 10003, ²Department of Psychological and Brain Sciences, Dartmouth College, Hanover, New Hampshire 03755, and ³Department of Psychology, Harvard University, Cambridge, Massachusetts 02138

Previous research shows that the amygdala automatically responds to a face's trustworthiness when a face is clearly visible. However, it is unclear whether the amygdala could evaluate such high-level facial information without a face being consciously perceived. Using a backward masking paradigm, we demonstrate in two functional neuroimaging experiments that the human amygdala is sensitive to subliminal variation in facial trustworthiness. Regions in the amygdala tracked how untrustworthy a face appeared (i.e., negative-linear responses) as well as the overall strength of a face's trustworthiness signal (i.e., nonlinear responses), despite faces not being subjectively seen. This tracking was robust across blocked and event-related designs and both real and computer-generated faces. The findings demonstrate that the amygdala can be influenced by even high-level facial information before that information is consciously perceived, suggesting that the amygdala's processing of social cues in the absence of awareness may be more extensive than previously described.

Key words: amygdala; backward masking; face; fMRI; social cues

Introduction

With only a glance, humans instantly form impressions of another's face. Such impressions occur spontaneously and are often beyond our conscious control (Zebrowitz and Montepare, 2008). They help us distinguish friend from foe, or those whom we should trust from those of whom we should be wary. Indeed, a mere 50 ms exposure to a face permits trait inferences that are highly correlated among multiple perceivers, indicating that facial cues provide reliable signals about another's underlying disposition (Bar et al., 2006; Willis and Todorov, 2006).

Previous behavioral studies suggest that face-based evaluations are underpinned by two fundamental dimensions, trustworthiness and dominance (Oosterhof and Todorov, 2008). Facial trustworthiness in particular accounts for the bulk of variance in social evaluation (Oosterhof and Todorov, 2008), and recent behavioral studies have provided preliminary evidence that individuals might be sensitive to trustworthiness without perceptual awareness (Todorov et al., 2009; Stewart et al., 2012). However, the neural basis of evaluating high-level social information such as trustworthiness from a face outside awareness remains unexplored.

Evaluations of trustworthiness reflect more general face valence and show correlated activity in the amygdala (Winston et al., 2002; Engell et al., 2007), a subcortical region involved in

processing the affective significance of social stimuli and important for a variety of social and emotional behaviors (Phelps and LeDoux, 2005; Adolphs, 2010). Consistent with the sensitivity of the amygdala to negative, threat-related stimuli, initial studies reported linear effects, with amygdala activation increasing for faces appearing less trustworthy (Winston et al., 2002; Engell et al., 2007). This effect held true regardless of task demands (Engell et al., 2007; Todorov and Engell, 2008; Todorov et al., 2011), suggesting that the amygdala may code trustworthiness implicitly when a face is clearly visible. More recent studies have reported quadratic effects, with amygdala activation increasing for faces appearing either more or less trustworthy, relative to neutral (Said et al., 2010; Todorov et al., 2011), potentially reflecting the coding of the salience or motivational relevance of a stimulus derived from a face's trustworthiness (Todorov et al., 2013). A recent meta-analysis found that both linear and nonlinear responses to trustworthiness coexist in different amygdala subregions (Mende-Siedlecki et al., 2013). However, it is currently unknown whether the amygdala is sensitive to trustworthiness before a face can reach perceptual awareness.

Just as the detection of threat-relevant stimuli by the amygdala is functionally adaptive (Öhman, 2005), so too might be its evaluation of a face's trustworthiness. Indeed, a large percentage of death throughout human history is a result of tribal conflicts and coalitional aggression, estimated to have had a substantial impact on human evolution (Chagnon, 1988). Thus, automatic evaluation of another's likelihood to harm or help via facial trustworthiness would facilitate survival and resource maintenance (McDonald et al., 2012). By facilitating the coding of another's trustworthiness before awareness, the amygdala could modulate cortical processes and motivate appropriate behavioral responses. Here, we describe two functional neuroimaging experiments designed to test whether the amygdala can respond to a

Received Dec. 4, 2013; revised June 26, 2014; accepted June 27, 2014.

Author contributions: J.B.F. designed research; R.M.S., Z.A.I., and E.A.H. performed research; J.B.F. and R.M.S. analyzed data; J.B.F. wrote the paper.

We thank Paul Whalen and Todd Heatherton for helpful comments.

The authors declare no competing financial interests.

Correspondence should be addressed to Dr. Jonathan B. Freeman, Department of Psychology, New York University, 6 Washington Place, New York, NY 10003. E-mail: jon.freeman@nyu.edu.

DOI:10.1523/JNEUROSCI.5063-13.2014

Copyright © 2014 the authors 0270-6474/14/3410573-09\$15.00/0

face's trustworthiness without perceptual awareness and to characterize the nature of this responsivity.

Materials and Methods

In Experiment 1, whole-brain fMRI data were collected during a backward masking paradigm involving three levels of masked facial trustworthiness (low, average, high), adapted from previous studies using a blocked design to maximize statistical power (Whalen et al., 1998; Kim et al., 2010). In Experiment 2, we extended the backward masking paradigm to a rapid event-related design that allowed us to test amygdala responsivity to a wider and fully continuous range of facial trustworthiness, and to directly compare neural activity between subliminal and supraliminal presentations.

Subjects. Twenty-one volunteers (16 females) between the ages of 18 and 22 years [mean (M) = 18.75 years] participated in Experiment 1, and 16 volunteers (12 females) between the ages of 18 and 35 years (M = 21.80 years) participated in Experiment 2. Two subjects in Experiment 1 and one subject in Experiment 2 were excluded because of reported awareness of the subliminal stimuli. All subjects in Experiments 1 and 2 were right-handed, native English speakers, and had normal or corrected-to-normal vision and no history of neurological disorders or use of psychoactive medications. They received partial course credit or monetary compensation.

Face stimuli. In a pretest, raters ($N = 10$) judged the trustworthiness of 300 neutral-affect, male and female faces from the Glasgow Unfamiliar Face Database (GUFDB; Burton et al., 2010), normalized for size and luminance, in randomized order on a 7-point Likert scale. Judgments were highly consistent across the raters (intraclass correlation coefficient = 0.912). For computer-generated targets, we used well validated faces developed by previous research using 3D face modeling to generate faces varying in trustworthiness (Oosterhof and Todorov, 2008). There were 24 computer-generated faces in each of the 3 conditions (low, average, high), corresponding to -2 SD (low), M (average), and $+2$ SD (high) of the trustworthiness dimension developed in previous work (Oosterhof and Todorov, 2008). All facial targets in Experiments 1 and 2 were emotionally neutral. See Figure 1 for task design and sample stimuli.

In Experiment 1, half of the functional runs presented real facial targets from the GUFDB, whereas the other half of functional runs presented computer-generated facial targets. For real targets, the pretest ratings were used to create 3 conditions: 24 low-trustworthy (M = 2.74, SE = 0.20), 24 average-trustworthy (M = 3.76, SE = 0.22), and 24 high-trustworthy (M = 4.67, SE = 0.21) faces. A mean rating per face, averaged across pretest raters, was computed. Low-trustworthy faces were rated as significantly less trustworthy than average-trustworthy faces ($t_{(46)} = 23.02, p < 0.0001$), which were rated significantly less trustworthy than high-trustworthy faces ($t_{(46)} = 19.02, p < 0.0001$). In Experiment 2, a total of 160 real male and female faces were used as the target stimuli, which were sampled evenly across the total distribution of 300 real faces varying in trustworthiness from Experiment 1 (stimuli from the GUFDB).

In functional runs involving real faces, neutral-affect Ekman faces were used as mask stimuli. In functional runs involving computer-generated faces, separate computer-generated faces that were neutral-affect, invariant on trustworthiness, and developed by the same previous research as the target stimuli were used as mask stimuli.

Design of Experiment 1. Experiment 1 used a backward masking paradigm involving 3 levels of masked facial trustworthiness (low, average, high), adapted from previous backward-masking studies using a blocked design to maximize statistical power (Whalen et al., 1998; Kim et al.,

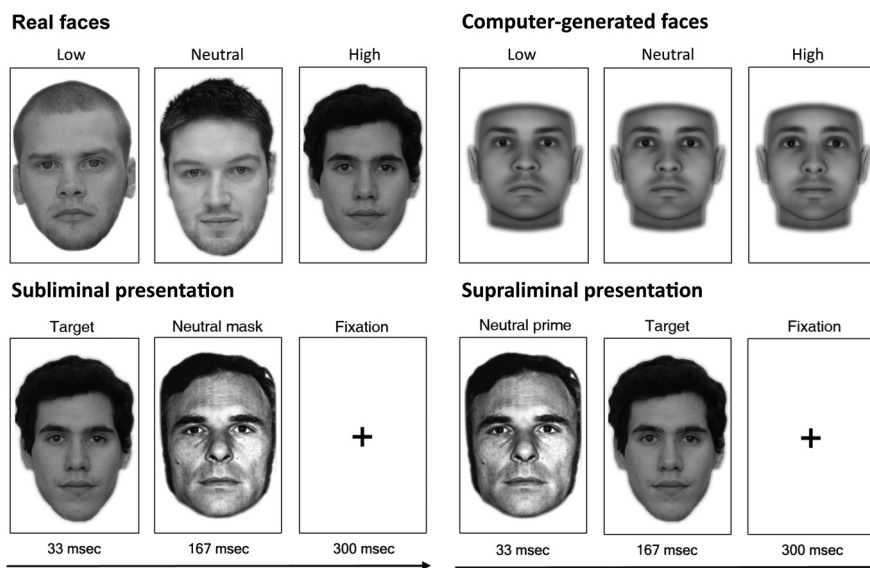


Figure 1. fMRI experimental procedure. Sample real stimuli (Experiments 1 and 2) and computer-generated stimuli (Experiment 1 only). For subliminal presentations (Experiments 1 and 2), the target face was presented for 33 ms and replaced by a neutral face mask for 167 ms that disrupted further visual processing of the target. This was followed by a 300 ms fixation period. For supraliminal presentations (Experiment 2), the target and neutral face were reversed.

2010). During fMRI, subjects viewed blocks of either low-, average-, or high-trustworthy facial targets that were masked by neutral face presentations. The targets in half of the functional scan runs were real faces from the GUFDB, described above, which were pre-rated on trustworthiness, obtained using the same camera and lighting, and not differing in low-level visual properties. Targets in the other half of the scan runs were computer-generated faces developed in previous studies using statistical face modeling to convey either low, average, or high trustworthiness (Oosterhof and Todorov, 2008). The computer-generated faces varied only in trustworthiness cues; all other perceptual information was controlled.

In the scanner, subjects passively viewed blocks of the target stimuli and masks across 10 functional runs. The first five runs consisted of real facial targets, and the last five runs consisted of computer-generated facial targets. During each run, subjects viewed 9 blocks. The order of block sequences was counterbalanced across subjects. Target trustworthiness varied between blocks, with 3 blocks per trustworthiness condition (low, average, and high trustworthiness). Each block consisted of 24 target faces unique to that block, whose order was randomized within the block. Five masks were randomly paired with targets per block. The targets were each centered on a black background in isolation for a duration of 33 ms. A neutral face mask then immediately replaced the prime for 167 ms, after which an interstimulus interval of 300 ms (fixation cross) ensued (Fig. 1). To maintain subjects' visual attention, subjects participated in a 1-back task in which they were asked to press a button if the same face (mask) was presented twice. The blocks were interleaved by 12 s fixation-cross baseline blocks in each run. Following the last block, a 10 s fixation-cross period completed the run.

Analysis of Experiment 1. Individual subjects' BOLD signals in Experiment 1 were modeled using a general linear model (GLM) with 6 predictors: 3 (low-trustworthy, average-trustworthy, and high-trustworthy faces) \times 2 (real faces, computer-generated faces). All predictors were modeled as boxcar functions across block durations and convolved with a two-gamma hemodynamic response function (HRF). First-level GLM analyses conducted on individual subjects' BOLD responses were submitted to a second-level random-effects analysis, treating subjects as a random factor.

Given an a priori hypothesis of amygdala involvement, we defined an anatomical region of interest (ROI) of the bilateral amygdala using the meta-analytic Neurosynth database (Yarkoni et al., 2011) as 10 mm spheres centered on $[\pm 18, -6, -11]$. Parameter estimates (β values)

were extracted from the bilateral amygdala ROI and submitted to a 3 (trustworthiness: low, average, high trustworthiness) \times 2 (stimulus type: real, computer-generated) repeated-measures ANOVA.

To further explore any additional effects of trustworthiness, we conducted random-effects analyses within a mask of the bilateral amygdala. Multiple statistical testing of voxels was corrected (false-positive rate $p < 0.05$) using a voxelwise threshold of $p < 0.05$ and a minimum cluster-size extent (k) of 421 mm³. The minimum cluster-size extent needed to maintain an experiment-wide α of 0.05 was empirically determined by a Monte Carlo simulation, accounting for spatial correlations between neighboring voxels (Forman et al., 1995).

Design of Experiment 2. In Experiment 2, we extended the backward masking paradigm to a rapid event-related design that allowed us to test amygdala responsiveness to a wider and fully continuous range of facial trustworthiness, and to directly compare neural activity between subliminal and supraliminal presentations. During fMRI, subjects took part in a modified event-related version of the backward masking paradigm using a considerably larger set of 160 target faces. Targets' trustworthiness varied trial-by-trial (each 2000 ms) under both subliminal and supraliminal conditions. Trustworthiness ratings were obtained for all faces postscan. Subjects passively viewed the target stimuli and neutral Ekman face masks across 4 functional runs. The first 2 runs presented targets subliminally, similar to Experiment 1. Targets were centered on a black background for 33 ms, which was immediately replaced by a neutral Ekman face mask for 167 ms. The mask was then followed by a fixation cross for 300 ms (Fig. 1). This 500 ms sequence repeated four consecutive times to maximize BOLD sensitivity to the targets and cover the length of a single TR (2000 ms). The second 2 runs involved supraliminal presentations, wherein the target faces and masks were reversed. Thus, a neutral Ekman face was centered on a black background for 33 ms, to be immediately followed by a target face for 167 ms and finally a 300 ms fixation cross. Like the first 2 subliminal runs, this sequence repeated four consecutive times (2000 ms). The design therefore ensured that visual information was identical across the subliminal versus supraliminal conditions. We refer to a single 2000 ms sequence as a "subliminal event" or "supraliminal event," respectively.

Of the total 160 faces used, two sets of 80 faces with matched trustworthiness distributions were then created. Each set was divided into 4 levels of 20 faces (low, medium-low, medium-high, high) to generate a number of presentation orders within a run that maximized statistical power for detecting parametric effects (i.e., to spread the trustworthiness variability across the run). For the first subliminal run, the 80 faces from one of the two sets (each face repeated twice) and an additional 20 baseline events (2000 ms of fixation cross) were used; for the second subliminal run, the 80 faces of the remaining set (each face repeated twice) and an additional 20 baseline events were used. The 2 supraliminal runs were identical to the 2 subliminal runs with the exception of reversing the target and mask stimuli. All events within runs were sequenced in a manner to optimize the efficiency of event-related BOLD signal estimation (Dale, 1999). Which face set was presented first versus second and the presentation order within runs were counterbalanced across subjects. A 6 s fixation-cross period divided the first and second half of each run (serving as a break), and another 6 s fixation-cross period completed each run. Note that it was not possible to counterbalance the order of subliminal versus supraliminal runs, as having supraliminal runs precede subliminal runs would heighten subjects' awareness of the target stimuli before their subliminal presentation.

After the scan, subjects were presented with each of the target faces one at a time in randomized order and rated their trustworthiness from 1 ("not at all") to 7 ("very much") using the keyboard. These ratings were used on a subject-by-subject, face-by-face basis to model BOLD responses.

Analysis of Experiment 2. In Experiment 2, subjects' postscan trustworthiness ratings were z -normalized. Individual subjects' BOLD signals were modeled using a GLM with 6 total predictors: 2 dichotomous predictors for modeling the presentation of subliminal and supraliminal events, 2 parametric predictors modeling a subject's postscan trustworthiness rating of the target for subliminal and supraliminal events (linear effects), and 2 parametric predictors modeling the square of a subject's

trustworthiness rating of the target for subliminal and supraliminal events (quadratic effects). All predictors were modeled as boxcar functions (the amplitude of which was parametrically varied, for parametric predictors) and convolved with a two-gamma HRF. First-level GLM analyses conducted on individual subjects' fMRI signal were submitted to a second-level random-effects analysis, treating subjects as a random factor.

Using the anatomical ROI of the bilateral amygdala (see details in Analysis of Experiment 1), β values were extracted to test for linear and quadratic trustworthiness effects. To identify clusters of activation within the amygdala exhibiting significant parametric effects, we conducted random-effects parametric analyses using a restricted mask of the bilateral amygdala. We controlled for multiple statistical testing of voxels within the bilateral amygdala mask (false-positive rate $p < 0.05$) using the same correction technique as in Experiment 1.

To better specify the nature of the quadratic modulation in the amygdala, β values associated with the individual 160 face stimuli were extracted for each subject. These were submitted to polynomial regression analysis using a multilevel generalized estimating equation (GEE) approach that can incorporate such nested, trial-by-trial data while accounting for the intracorrelations due to repeated measurements (Zeger and Liang, 1986). In our case, this included intracorrelations associated with individual subjects and with individual face stimuli. Separately for the left and right amygdala, β values associated with the 160 targets were regressed onto linear and quadratic components of subjects' postscan trustworthiness ratings.

Postscan discrimination tasks. To provide an objective measure of awareness in Experiments 1 and 2, we used postscan discrimination tasks. While still in the scanner, subjects were presented with facial targets one at a time in randomized order using a procedure virtually identical to the one used in the main fMRI tasks. Before starting the discrimination task, subjects were informed of target face presentations occurring before the masks, and were asked to decide their gender. The accuracy of gender categorizations served as the measure of objective awareness. Each trial began with a 500 ms fixation cross, followed by the target face. After 33 ms, the target was immediately replaced by the mask for 167 ms. In the Experiment 1 discrimination task, the mask was followed by a prompt for gender categorization (male or female?). In the Experiment 2 discrimination task, the fixation cross, target face, and mask repeated four consecutive times (as was done in the Experiment 2 fMRI task) before the categorization prompt appeared. Subjects were instructed to categorize the gender of the target as quickly and accurately as possible using a button press; they were given unlimited time to make their response. Differing from the main fMRI tasks, trials were self-paced. Subjects completed several practice trials before starting the task. In the Experiment 1 discrimination task, subjects completed 144 trials (72 real face targets and 72 computer-generated face targets). Because Experiment 1 involved only male targets, we used a representative half of the faces used in the main fMRI task in addition to an equal number of filler female faces matched for visual properties. In the Experiment 2 discrimination task, subjects completed 160 trials (the 160 real faces—half male, half female—used in the main fMRI task). The same Ekman and computer-generated mask stimuli were used for real and computer-generated face targets, respectively.

Discrimination of facial trustworthiness. Although our postscan discrimination tasks in Experiments 1 and 2 assessed perceptions of gender to ensure a lack of discriminability for masked targets' facial characteristics in general, trustworthiness could potentially have remained discriminable despite gender being indiscriminable with our masking procedure. To alleviate this concern, an additional behavioral experiment was conducted in which 16 volunteers (10 females) between the ages of 18 and 24 years ($M = 19.19$ years) participated. The procedure was identical to the postscan discrimination tasks of Experiments 1 and 2, except that participants were prompted to judge the target's trustworthiness (untrustworthy or trustworthy?) rather than gender. Trials were self-paced and subjects were given unlimited time to make their response. After completing several practice trials, subjects were presented with 211 trials, comprising the 72 non-filler targets used in the Experiment 1 discrimination task and 160 targets used in the Experiment 2




Figure 2. Stronger bilateral amygdala activation to low-trustworthy faces. Coronal slice ($y = -5$), depicting stronger responses to low-trustworthy targets versus average- or high-trustworthy targets from a random-effects analysis targeting the bilateral amygdala ($p < 0.05$, corrected). Bar plots depict mean β values for the 3 block types. Error bars indicate SEM.

discrimination task (there were 21 targets that overlapped in both experiments, resulting in a total of 211 trials). Following this initial phase of the experiment, subjects viewed all targets again in a new randomized order, one at a time, with unlimited exposure (no masking), and they were asked to judge trustworthiness along a 6-point Likert scale. We used a 6-point rather than a 7-point scale to be able to dichotomize judgments (0–3 = untrustworthy; 4–6 = trustworthy), thereby permitting signal detection analysis and a controlling of response bias.

fMRI acquisition and preprocessing. In both experiments, subjects were scanned using a 3T Philips Intera Achieva Scanner (Philips Medical Systems) equipped with a SENSE birdcage head coil in the Dartmouth Brain Imaging Center. All stimuli were back-projected onto a screen visible via a mirror mounted on the MRI head coil (visual angle $\sim 13.5 \times 13.5^\circ$). Anatomical images were acquired using a T1-weighted protocol (256×256 matrix, 128 1.33 mm transverse slices). Functional images were acquired using a single-shot gradient echo EPI sequence (TR = 2000 ms, TE = 35 ms). Thirty-five interleaved oblique-axial slices ($3 \text{ mm} \times 3 \text{ mm} \times 4 \text{ mm}$ voxels; no slice gap) parallel to the AC-PC line were obtained. Analysis of the imaging data was conducted using BrainVoyagerQX (Brain Innovation). Functional imaging data preprocessing included 3D motion correction, slice-timing correction (sinc interpolation), spatial smoothing using a 3D Gaussian filter (8 mm FWHM), and voxelwise linear detrending and high-pass filtering of frequencies (>3 cycles per time course). Structural and functional data of each subject were transformed to standard Talairach stereotaxic space.

Results

Experiment 1

Experiment 1 presented subjects with subliminally presented faces that varied on three levels of trustworthiness (low, average, high). Subjects were questioned about their subjective awareness of the subliminal stimuli after the scan; two subjects reported subjective awareness and were excluded. Postscan discrimination data from the remaining subjects were analyzed using signal detection to appropriately control for response bias; signal was arbitrarily defined as female. Thus, perceptual discriminability (d') was computed as the percentage of masked female primes that were successfully categorized as female (hits), adjusted for the percentage of male masked primes that were erroneously categorized as female (false alarms): $d' = z\text{-score}(\% \text{ hits}) - z\text{-score}(\% \text{ false alarms})$, with chance performance = 0 ± 1.74 . No included subject's discriminability (d') rose significantly above chance (all p values < 1.74), and d' overall was quite low ($M = 0.17$, $SE = 0.11$), ensuring that the masked stimuli were below subjects' awareness.

Given an a priori hypothesis of amygdala responsivity, we examined activation in the bilateral amygdala using anatomically

defined, 10 mm spherical ROIs. A repeated-measures ANOVA conducted on parameter estimates (β values) extracted from the bilateral amygdala ROI revealed a significant effect of facial trustworthiness, $F_{(2,36)} = 3.80$, $p = 0.032$. Paired-samples t tests indicated that this effect was driven specifically by low levels of trustworthiness. Low-trustworthy targets elicited stronger activation than average-trustworthy targets ($t_{(18)} = 2.12$, $p = 0.048$); however, average- and high-trustworthy targets were not distinguished ($t_{(18)} = 0.54$, $p = 0.593$; Fig. 2B). The main effect of stimulus type (real vs computer-generated; $F_{(1,18)} = 0.40$, $p = 0.534$) and the interaction ($F_{(2,36)} = 0.33$, $p = 0.720$) were not significant; further analyses therefore were collapsed across real and computer-generated facial targets. To ensure that any possible weak, albeit nonsignificant, discriminability (d') of the masked stimuli was not driving this effect, subjects' d' was included as a covariate. However, responsivity in the bilateral amygdala to trustworthiness persisted when controlling for d' ($F_{(2,34)} = 3.477$, $p = 0.042$), indicating that discriminability had a negligible impact on the effect. Furthermore, the interaction between trustworthiness and d' was not significant ($F_{(2,34)} = 0.069$, $p = 0.933$), alleviating the concern that any potential visibility confounded the amygdala responses to trustworthiness.

To examine the possibility of additional trustworthiness effects, a random-effects ANOVA tested for clusters of activation that exhibited modulation by trustworthiness using a restricted mask of the bilateral amygdala. This revealed clusters in the right ($x = 22$, $y = -4$, $z = -10$, 429 mm^3 , $F = 4.11$) and left ($x = -21$, $y = -9$, $z = -9$, 769 mm^3 , $F = 3.81$) amygdala that showed especially strong responses to low-trustworthy targets ($p < 0.05$, corrected; Fig. 2), corroborating the anatomical ROI analysis.

To specifically identify any amygdala regions that were responsive to both low and high trustworthiness (a nonlinear effect), we conducted a random-effects contrast of [low + high] $>$ average. This revealed a cluster in the right amygdala located more posteriorly ($x = 21$, $y = -12$, $z = -10$, 677 mm^3 , $F = 2.32$; $p < 0.05$, corrected). The β values were extracted from this region, and a within-subject ANOVA contrast [1, -2, 1] confirmed a quadratic response, $F_{(1,18)} = 9.380$, $p = 0.007$. Planned comparisons indicated that, relative to average-trustworthy faces, both low-trustworthy ($t_{(18)} = 2.70$, $p = 0.015$) and high-trustworthy ($t_{(18)} = 2.11$, $p = 0.049$) faces more strongly engaged this region (one-tailed tests for directional hypothesis; Fig. 3). Similar to the negative-linear effect in the bilateral amygdala, the

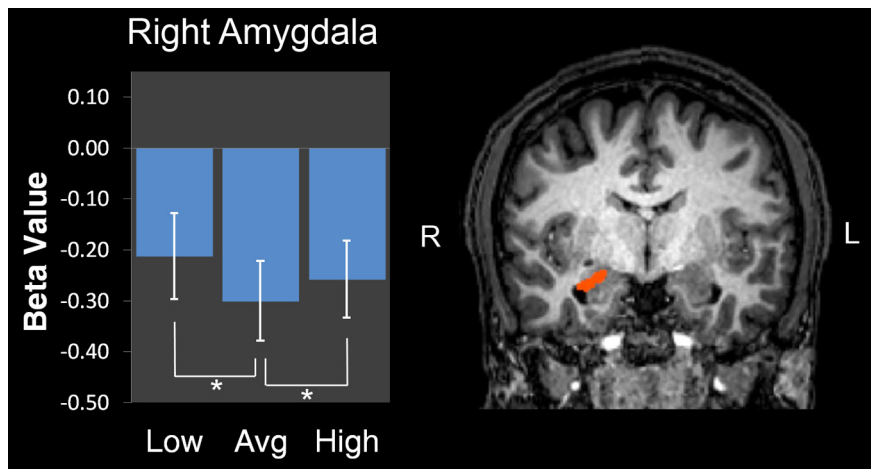


Figure 3. Region of the right amygdala responsive to both low and high levels of trustworthiness. Coronal slice ($y = -7$) depicting a posterior region of the right amygdala that responded more strongly to low- and high-trustworthy faces than average-trustworthy faces from a random-effects analysis targeting the bilateral amygdala ($p < 0.05$, corrected). Bar plots depict mean β values for the 3 block types. Error bars indicate SEM.

quadratic effect in this posterior region of the right amygdala remained significant after including d' as a covariate ($F_{(1,17)} = 7.642$, $p = 0.013$); further, the interaction between trustworthiness and d' was not significant ($F_{(1,17)} = 0.002$, $p = 0.964$), thus ensuring that any potential visibility did not confound this result.

Given the critical role of the fusiform cortex in face perception (Haxby et al., 2000), exploratory random-effects analyses tested for such linear and nonlinear effects of trustworthiness within a mask of the bilateral fusiform cortex; however, no effects were present that survived correction.

Thus, regions in the bilateral amygdala exhibited especially strong activation for low-trustworthy faces when presented subliminally. Moreover, a region of the posterior right amygdala exhibited a nonlinear pattern of response, showing stronger activation to both low and high trustworthiness. These findings extend previous reports of linear and nonlinear amygdala responses to trustworthiness during supraliminal presentations to the subliminal level.

Interestingly, amygdala activation was reduced below baseline across the three trustworthiness conditions. This may be because of deactivations of the amygdala commonly observed in tasks required high-level cognitive processing. Specifically, because experimental blocks involved a 1-back task used to maintain subjects' attention, the demands of the 1-back task may have overall reduced amygdala responses relative to baseline blocks that required the mere passive viewing of a fixation cross. Previous work suggests that cognitively demanding tasks (e.g., n -back tasks) suppress amygdala response to visual stimuli (Drevets and Raichle, 1998), even when those stimuli are task-relevant (Kellermann et al., 2012). Thus, one explanation for reduced amygdala responses during experimental blocks is the additional cognitive demands associated with those blocks. Despite an overall reduced response in the amygdala, however, the primary result and the amygdala's relative response to subliminal variation in facial trustworthiness is clear.

In Experiment 2, we expanded on these results by testing amygdala responsivity to a full continuum of trustworthiness using a rapid event-related design, and directly comparing with supraliminal presentations.

Experiment 2

After the scan, one subject reported subjective awareness of the targets and was excluded. As in Experiment 1, postscan forced-choice discrimination data were analyzed using signal detection. No subject's discriminability (d') rose significantly above chance (all p values < 1.74), and overall d' was even negative ($M = -0.09$, $SE = 0.06$), ensuring that the masked stimuli were below subjects' awareness.

Parametric analyses simultaneously modeled linear and nonlinear (quadratic) predictors based on subjects' postscan trustworthiness ratings. Activation averaged across the anatomical ROI of the bilateral amygdala was significantly modulated by a quadratic trustworthiness effect ($t_{(14)} = 2.25$, $p = 0.041$); no linear effect emerged ($t_{(14)} = 0.30$, $p = 0.769$). To investigate further, a random-effects parametric analysis targeting the bilateral

amygdala was probed for clusters exhibiting significant quadratic responses ($p < 0.05$, corrected). This revealed clusters in the left ($x = -22$, $y = -5$, $z = -17$; $k = 2495$ mm³; $t = 3.64$) and right ($x = 19$, $y = -6$, $z = -16$; $k = 1559$ mm³; $t = 3.51$) amygdala (Fig. 4, top), consistent with the anatomical ROI analysis. The β values were extracted to examine the parametric effects separately for subliminal and supraliminal conditions. The quadratic effect was significant during supraliminal presentations (left: $t_{(14)} = 2.17$, $p = 0.048$; right: $t_{(14)} = 2.55$, $p = 0.023$). Critically, it was also significant during subliminal presentations (left: $t_{(14)} = 2.17$, $p = 0.048$; right: $t_{(14)} = 2.35$, $p = 0.034$). Linear effects were not significant in these clusters for either supraliminal (left: $t_{(14)} = -0.15$, $p = 0.885$; right: $t_{(14)} = -1.19$, $p = 0.252$) or subliminal (left: $t_{(14)} = 1.81$, $p = 0.091$; right: $t_{(14)} = 0.25$, $p = 0.804$) conditions (Fig. 4, middle). A random-effects parametric analysis testing for significant linear responses revealed no clusters of activation that survived correction.

To better specify the nature of the quadratic modulation, an additional model was constructed to extract β values uniquely associated with the individual 160 face stimuli. Extracted β values were submitted to polynomial regression analyses using a GEE approach that can incorporate such nested, trial-by-trial data while accounting for intracorrelations due to repeated measurements (Zeger and Liang, 1986). In both the left and right amygdala, the linear effect of trustworthiness was not significant (left: $B = 0.010$, $Z = 0.38$, $p = 0.705$; right: $B = 0.002$, $Z = 0.06$, $p = 0.953$), whereas the quadratic effect was highly significant (left: $B = 0.040$, $Z = 2.66$, $p = 0.008$; right: $B = 0.037$, $Z = 2.44$, $p = 0.015$). This quadratic effect did not interact with presentation condition (subliminal vs supraliminal; left: $B = 0.0166$, $Z = 0.61$, $p = 0.540$; right: $B = 0.0319$, $Z = 1.12$, $p = 0.262$), and thus, further analyses collapsed across them. Plotting mean β values in the left and right amygdala revealed a U-shaped encoding function. For both supraliminal and subliminal presentations, the amygdala increased in activation as faces approached the low and high ends of the continuum, relative to average faces at the middle (Fig. 4, middle and bottom). As in Experiment 1, d' was included as a covariate to control for any possible influence of weak discriminability of masked stimuli. The quadratic effect remained significant after including d' as a covariate (left: $B =$

0.0416, $Z = 2.77$, $p = 0.006$; right: $B = 0.0402$, $Z = 2.66$, $p = 0.008$) and did not significantly interact with d' (left: $B = -0.0734$, $Z = -0.92$, $p = 0.358$; right: $B = -0.1122$, $Z = -1.43$, $p = 0.152$), eliminating concerns of possible visibility confounding the effect.

Finally, as in Experiment 1, exploratory random-effects analyses tested for any linear and nonlinear responses within a mask of the bilateral fusiform cortex; however, no effects were present that survived correction.

Thus, the results of Experiment 2 replicate the nonlinear effects of trustworthiness obtained in previous studies using supraliminal presentations, and show that such effects extend to subliminal exposures outside subjects' awareness.

We should note that, because the amygdala can exhibit adaptation to emotionally significant stimuli in masking paradigms (Whalen et al., 1998), it is possible that the visible repetition of a small number of Ekman masks in the subliminal condition may have elicited adaptation effects not present in the supraliminal condition (which visibly presented 160 targets without repetition). This may be additionally important because amygdala responses can be shaped by the conjunction of masked and masking stimuli (Kim et al., 2010). However, because the nonlinear amygdala response to trustworthiness was statistically identical across the subliminal and supraliminal conditions, this is unlikely to have been an issue in the present experiment.

Discrimination of masked trustworthiness

An additional behavioral experiment was run to ensure that facial trustworthiness in particular was not discriminable from the masking procedure we used (see Materials and Methods). No subject reported subjective awareness of the targets after the experiment. Forced-choice discrimination data (0 = untrustworthy; 1 = trustworthy) were analyzed using signal detection. A face's "correct" trustworthiness was defined by a subject's own rating of that specific face following the discrimination task, in which a 6-point Likert scale was used to be able to dichotomize ratings (1–3 = untrustworthy; 4–6 = trustworthy). No subject's d' rose significantly above chance (all p values < 1.74), and overall d' was quite small ($M = 0.03$, $SE = 0.06$), ensuring that the trustworthiness of masked faces was not discriminable.

Correspondingly, we should note that although subjects did not report any subjective awareness of masked faces, it is possible a more objective measure of awareness would have shown evidence of subjects detecting some masked stimuli as faces (as opposed to non-faces). Indeed, the basic detection of a face is dissociable from the discrimination of facial characteristics such as trustworthiness or gender. For instance, one recent study found evidence for accurate discrimination of facial emotion de-

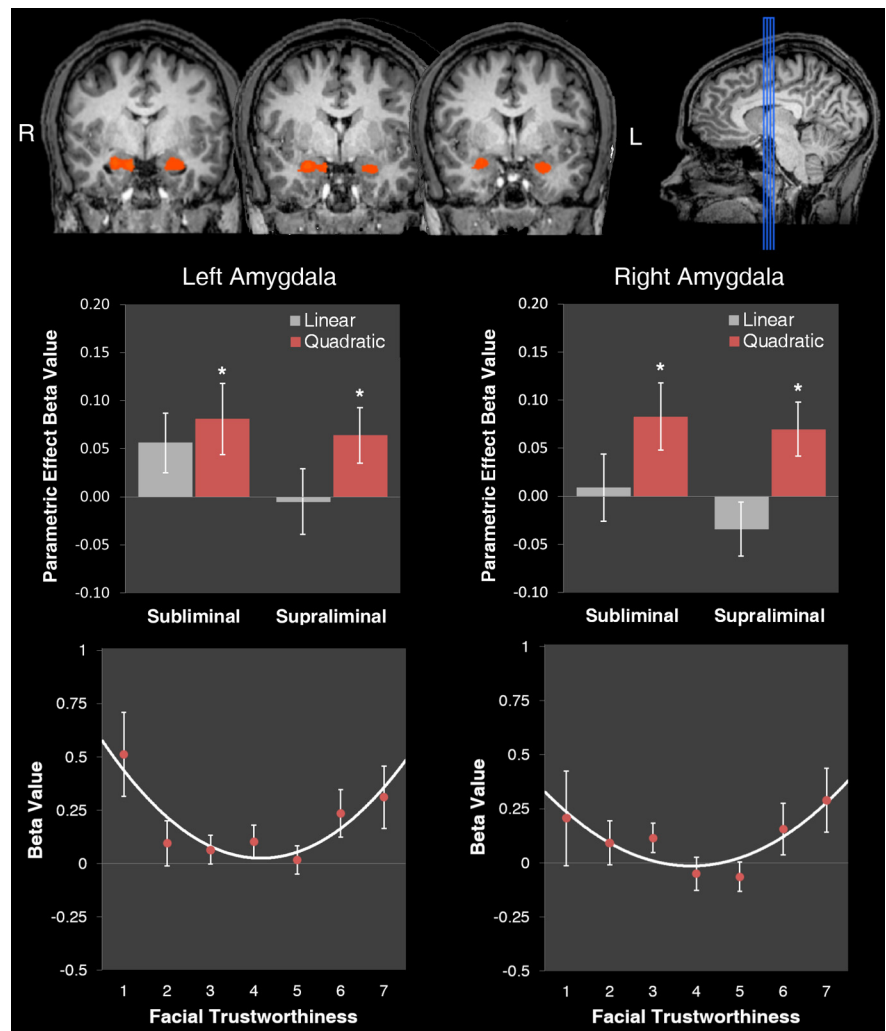


Figure 4. Nonlinear amygdala responses to facial trustworthiness during subliminal and supraliminal presentations. Top, Coronal slices ($y = -1 \pm 2$) depicting responses associated with a quadratic trustworthiness effect from a random-effects analysis targeting the bilateral amygdala ($p < 0.05$, corrected). Middle, Mean β values for the linear and quadratic parametric predictors, indexing the strength of the correlations between amygdala response and linear vs quadratic components of trustworthiness, separately for left amygdala (left) and right amygdala (right). Bottom, Mean β values for amygdala BOLD response as a function of facial trustworthiness (for both supraliminal and subliminal conditions), separately for left amygdala (left) and right amygdala (right). Error bars indicate SEM.

spite the facial stimuli being so impoverished they were explicitly categorized as shoes (Seirafi et al., 2013). Nevertheless, what is critical is that facial trustworthiness in particular was not discriminable in the subliminal presentations we used, as indicated by the aforementioned analysis. This specifically shows that the amygdala can track high-level social information from a face (e.g., trustworthiness) without that information being consciously perceived.

Discussion

Across two experiments, our findings demonstrate that the human amygdala is automatically responsive to a face's trustworthiness in the absence of perceptual awareness. The amygdala has been shown to code trustworthiness in an attention-independent fashion when faces are clearly visible (Engell et al., 2007; Todorov et al., 2011). However, the notion that a high-level facial characteristic such as trustworthiness may be assessed before a face is consciously perceived has only been recently suggested by behavioral studies (Todorov et al., 2009;

Stewart et al., 2012), and the role of the amygdala in this process has remained unclear.

The present results are striking in that they raise a conundrum as to how the amygdala could evaluate such complex social information under constricted processing. Researchers have documented amygdala responses to subliminal fear expressions (Whalen et al., 1998), which have been suggested by some to be related to an early threat-detection mechanism driven by a retino-collicular-pulvinar pathway that responds to salient, behaviorally relevant stimuli (Morris et al., 1999; Davis and Whalen, 2001; LeDoux, 2007). However, amygdala responses to subliminal fearful expressions may be driven by simple visual cues such as enlarged eye whites (Whalen et al., 2004). In contrast, the features associated with facial trustworthiness are considerably more complex (Oosterhof and Todorov, 2008), although at extreme levels, they do overlap with emotional expressions of anger and joy (Oosterhof and Todorov, 2009), and some studies have found amygdala activation to masked angry faces (Morris et al., 1998). Nonconscious processing of faces has been reported within ventral temporal cortex alone (Green et al., 2005; Sterzer et al., 2009), where detailed facial information can be extracted (Haxby et al., 2000; Kanwisher and Yovel, 2006). Limited processing via the ventral temporal cortex could potentially drive amygdala activation to subliminal stimuli as well, and thus there are multiple routes through which facial trustworthiness could affect amygdala responsivity without perceptual awareness (Pessoa and Adolphs, 2010; Tamietto and de Gelder, 2010). Such ventral-temporal processing during masked exposures has been shown to exhibit face-specific encoding (Sterzer et al., 2009), which could potentially contribute to amygdala responsivity (Pessoa and Adolphs, 2010; Tamietto and de Gelder, 2010). Indeed, facial trustworthiness is processed extensively by face-specific areas in the ventral-visual stream, such as the fusiform cortex (Said et al., 2010; Mende-Siedlecki et al., 2013; Todorov et al., 2013).

It is also quite possible that both cortically and subcortically driven routes to the amygdala could be involved in the subliminal effects, with the critical question being the degree of cortico-subcortical interaction. In this perspective, it is the feedback loops between subcortical and cortical areas that are thought to be critical for awareness, rather than the involvement of particular brain regions. Thus, masked exposures to facial trustworthiness may have been sufficient to modulate amygdala responses because of enough bottom-up information arriving through feedforward connections (whether they be subcortical or cortical), but not sufficient to engage the more elaborate reentrant cortico-subcortical feedback important for awareness (Williams et al., 2006; Pessoa and Adolphs, 2010). As such, masked trustworthiness could have activated limited processing in any number of regions that possibly contributed to amygdala responses here, but it may not have been able to be consciously perceived due to a lack of cortico-subcortical interaction. Future research could clarify the specific pathways that underlie the amygdala's tracking of trustworthiness outside awareness.

If, however, ventral-stream face processing participated in the amygdala responsivity here, one would have expected face-processing areas (e.g., fusiform cortex) to be modulated by trustworthiness. In the present experiments, the fusiform cortex was not reliably modulated by trustworthiness, consistent with some previous studies finding robust amygdala effects of trustworthiness without significant fusiform effects (Engell et al., 2007). One possibility is that our current methodological approach was not sensitive enough to detect trustworthiness effects in the fusiform

cortex. For instance, because the fusiform cortex is a critical face-processing region, it exhibits particularly strong adaptation to visual facial features (Kanwisher and Yovel, 2006). Because of the nature of our masking procedure and because the target and mask stimuli were both emotionally neutral faces and highly similar, it is possible that fusiform responses were confounded by adaptation effects between target and mask stimuli. If these effects occurred, they would have diminished sensitivity to trustworthiness responses in the fusiform cortex considerably. Future work might therefore consider using different masking techniques to focus on fusiform responsivity. There is also the possibility of a genuine dissociation between amygdala and fusiform modulation by trustworthiness, such that the amygdala exhibits responses to subliminally presented trustworthiness without the fusiform exhibiting such responses. If this were true, however, reliable fusiform modulation would have been observed during the supraliminal condition, which was not the case. Further studies that target fusiform responsivity directly could address these possibilities.

It should also be noted that although trustworthiness is regarded as a fundamental dimension of social evaluation (Oosterhof and Todorov, 2008), clearly, it is composed of lower-level featural dimensions that jointly contribute to its perception. For example, features on an emotionally neutral face exhibiting subtle structural overlap with particular emotion expressions (e.g., anger or joy) contribute to perceived trustworthiness (Said et al., 2009). Complex arrangements of such lower-level facial features together form this fundamental trustworthiness dimension, which is spontaneously perceived on encountering a face and accounts for the majority of variance in social evaluation (Engell et al., 2007; Stewart et al., 2012). Here we show that this dimension can sensitively modulate amygdala responses before reaching awareness.

It is only recently that researchers have begun to examine linear and nonlinear responses to facial trustworthiness in tandem (Said et al., 2010; Todorov et al., 2011; Mende-Siedlecki et al., 2013). In some studies, as in Experiment 1, such effects cohabit within different subregions of the amygdala (Todorov et al., 2013); in other studies, as in Experiment 2, quadratic effects trump linear effects entirely (Said et al., 2010). Work in non-human primates has also converged on the finding of both forms of coding in the amygdala, with one subregion showing linear responses to threatening faces (coding negative valence), and another showing nonlinear responses to both threatening and appealing faces relative to neutral (coding salience; Hoffman et al., 2007). In humans, neuroimaging work more generally has dissociated amygdala subregions coding for the valence versus salience of stimuli in ambiguous learning situations, e.g., a novel social encounter (Whalen, 1998; Kim et al., 2003; Davis et al., 2010). The results of Experiment 1 are thus consistent with this prior research, finding coexisting linear and nonlinear representations; however, Experiment 2 obtained evidence only for a nonlinear representation. Although also consistent with prior work, the difference could be driven by the nature of the blocked versus event-related tasks. Repeated presentations in the blocked design of Experiment 1 were similar to the task design of Kim et al. (2003) and may have induced a task context that increased the importance of tracking valence over salience. Alternatively, the event-related design of Experiment 2 used a wider, continuous range of trustworthiness, which may have increased sensitivity to nonlinear effects (Todorov et al., 2013). Despite these differences between blocked versus event-related presentations, both experiments provide clear support for our primary hypothesis that the

amygdala is automatically responsive to facial trustworthiness without perceptual awareness. Regions in the amygdala track how untrustworthy an unseen face appears as well as the overall strength of the trustworthiness signal. Moreover, this tracking generalizes across real and computer-generated faces, where trustworthiness was both measured and manipulated, respectively.

Whereas the negative-linear effects found here are consistent with the amygdala's role in vigilance for threats and tracking of valence (Davis and Whalen, 2001; LeDoux, 2007), nonlinear effects are consistent with the amygdala's processing of salience (Zald, 2003) and motivational relevance (Phelps and LeDoux, 2005; Cunningham et al., 2008; Adolphs, 2010). Faces with stronger cues for untrustworthiness or trustworthiness are more motivationally relevant, as these cues spontaneously elicit approach and avoidance behaviors (Slepian et al., 2012). Both forms of coding are consistent with the amygdala's importance for interpreting implicit social signals (Adolphs et al., 1998; Freeman et al., 2010; Heberlein and Adolphs, 2004). Faces that appear more untrustworthy and likely to inflict harm, or faces with a stronger trustworthiness signal in general, would benefit more from automatic amygdala responsivity, which could adaptively modulate cortical processes and motivate appropriate social behavior (Davis and Whalen, 2001; Phelps and LeDoux, 2005).

In summary, we demonstrated in two experiments that the amygdala is sensitive to subliminal variation in facial trustworthiness. Thus, the amygdala can be influenced by even high-level facial information before that information is consciously perceived. These findings provide evidence that the amygdala's processing of social cues in the absence of awareness may be more extensive than previously described.

References

- Adolphs R (2010) What does the amygdala contribute to social cognition? *Ann N Y Acad Sci* 1191:42–61. [CrossRef Medline](#)
- Adolphs R, Tranel D, Damasio AR (1998) The human amygdala in social judgment. *Nature* 393:470–474. [CrossRef Medline](#)
- Bar M, Neta M, Linz H (2006) Very first impressions. *Emotion* 6:269–278. [CrossRef Medline](#)
- Burton AM, White D, McNeill A (2010) The Glasgow face matching test. *Behav Res Methods* 42:286–291. [CrossRef Medline](#)
- Chagnon NA (1988) Life histories, blood revenge, and warfare in a tribal population. *Science* 239:985–992. [CrossRef Medline](#)
- Cunningham WA, Van Bavel JJ, Johnsen IR (2008) Affective flexibility: evaluative processing goals shape amygdala activity. *Psychol Sci* 19:152–160. [CrossRef Medline](#)
- Dale AM (1999) Optimal experimental design for event-related fMRI. *Hum Brain Mapp* 8:109–114. [CrossRef Medline](#)
- Davis FC, Johnstone T, Mazzulla EC, Oler JA, Whalen PJ (2010) Regional response differences across the human amygdaloid complex during social conditioning. *Cereb Cortex* 20:612–621. [CrossRef Medline](#)
- Davis M, Whalen PJ (2001) The amygdala: vigilance and emotion. *Mol Psychiatry* 6:13–34. [CrossRef Medline](#)
- Drevets WC, Raichle ME (1998) Reciprocal suppression of regional cerebral blood flow during emotional versus higher cognitive processes: implications for interactions between emotion and cognition. *Cognition Emotion* 12:353–385. [CrossRef](#)
- Engell AD, Haxby JV, Todorov A (2007) Implicit trustworthiness decisions: automatic coding of face properties in the human amygdala. *J Cogn Neurosci* 19:1508–1519. [CrossRef Medline](#)
- Forman SD, Cohen JD, Fitzgerald M, Eddy WF, Mintun MA, Noll DC (1995) Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): use of a cluster-size threshold. *Magn Reson Med* 33:636–647. [CrossRef Medline](#)
- Freeman JB, Schiller D, Rule NO, Ambady N (2010) The neural origins of superficial and individuated judgments about ingroup and outgroup members. *Hum Brain Mapp* 31:150–159. [CrossRef Medline](#)
- Green MF, Glahn D, Engel SA, Nuechterlein KH, Sabb F, Strojwas M, Cohen MS (2005) Regional brain activity associated with visual backward masking. *J Cogn Neurosci* 17:13–23. [CrossRef Medline](#)
- Haxby JV, Hoffman EA, Gobbini MI (2000) The distributed human neural system for face perception. *Trends Cogn Sci* 4:223–233. [CrossRef Medline](#)
- Heberlein AS, Adolphs R (2004) Impaired spontaneous anthropomorphizing despite intact perception and social knowledge. *Proc Natl Acad Sci U S A* 101:7487–7491. [CrossRef Medline](#)
- Hoffman KL, Gothard KM, Schmid MC, Logothetis NK (2007) Facial-expression and gaze-selective responses in the monkey amygdala. *Curr Biol* 17:766–772. [CrossRef Medline](#)
- Kanwisher N, Yovel G (2006) The fusiform face area: a cortical region specialized for the perception of faces. *Philos Trans R Soc Lond B Biol Sci* 361:2109–2128. [CrossRef Medline](#)
- Kellermann TS, Sternkopf MA, Schneider F, Habel U, Turetsky BI, Zilles K, Eickhoff SB (2012) Modulating the processing of emotional stimuli by cognitive demand. *Soc Cogn Affect Neurosci* 7:263–273. [CrossRef Medline](#)
- Kim H, Somerville LH, Johnstone T, Alexander AL, Whalen PJ (2003) Inverse amygdala and medial prefrontal cortex responses to surprised faces. *Neuroreport* 14:2317–2322. [CrossRef Medline](#)
- Kim MJ, Loucks RA, Neta M, Davis FC, Oler JA, Mazzulla EC, Whalen PJ (2010) Behind the mask: the influence of mask-type on amygdala response to fearful faces. *Soc Cogn Affect Neurosci* 5:363–368. [CrossRef Medline](#)
- LeDoux JE (2007) The amygdala. *Curr Biol* 17:R868–R874. [CrossRef Medline](#)
- McDonald MM, Navarrete CD, Van Vugt M (2012) Evolution and the psychology of intergroup conflict: the male warrior hypothesis. *Philos Trans R Soc B Biol Sci* 367:670–679. [CrossRef](#)
- Mende-Siedlecki P, Said CP, Todorov A (2013) The social evaluation of faces: a meta-analysis of functional neuroimaging studies. *Soc Cogn Affect Neurosci* 8:285–299. [CrossRef Medline](#)
- Morris JS, Ohman A, Dolan RJ (1998) Conscious and unconscious emotional learning in the human amygdala. *Nature* 393:467–470. [CrossRef Medline](#)
- Morris JS, Ohman A, Dolan RJ (1999) A subcortical pathway to the right amygdala mediating “unseen” fear. *Proc Natl Acad Sci U S A* 96:1680–1685. [CrossRef Medline](#)
- Öhman A (2005) The role of the amygdala in human fear: automatic detection of threat. *Psychoneuroendocrinology* 30:953–958. [CrossRef Medline](#)
- Oosterhof NN, Todorov A (2008) The functional basis of face evaluation. *Proc Natl Acad Sci U S A* 105:11087–11092. [CrossRef Medline](#)
- Oosterhof NN, Todorov A (2009) Shared perceptual basis of emotional expressions and trustworthiness impressions from faces. *Emotion* 9:128–133. [CrossRef Medline](#)
- Pessoa L, Adolphs R (2010) Emotion processing and the amygdala: from a ‘low road’ to ‘many roads’ of evaluating biological significance. *Nat Rev Neurosci* 11:773–783. [CrossRef Medline](#)
- Phelps EA, LeDoux JE (2005) Contributions of the amygdala to emotion processing: from animal models to human behavior. *Neuron* 48:175–187. [CrossRef Medline](#)
- Said CP, Sebe N, Todorov A (2009) Structural resemblance to emotional expressions predicts evaluation of emotionally neutral faces. *Emotion* 9:260–264. [CrossRef Medline](#)
- Said CP, Dotsch R, Todorov A (2010) The amygdala and FFA track both social and non-social face dimensions. *Neuropsychologia* 48:3596–3605. [CrossRef Medline](#)
- Seirafi M, De Weerd P, de Gelder BL (2013) Emotion categorization does not depend on explicit face categorization. *J Vis* 13(2):12–19. [CrossRef](#)
- Slepian ML, Young SG, Rule NO, Weisbuch M, Ambady N (2012) Embodied impression formation: social judgments and motor cues to approach and avoidance. *Soc Cogn* 30:232–240. [CrossRef](#)
- Sterzer P, Jalkanen L, Rees G (2009) Electromagnetic responses to invisible face stimuli during binocular suppression. *Neuroimage* 46:803–808. [CrossRef Medline](#)
- Stewart LH, Ajina S, Getov S, Bahrami B, Todorov A, Rees G (2012) Unconscious evaluation of faces on social dimensions. *J Exp Psychol Gen* 141:715–727. [CrossRef Medline](#)
- Tamietto M, de Gelder B (2010) Neural bases of the non-conscious perception of emotional signals. *Nat Rev Neurosci* 11:697–709. [CrossRef Medline](#)
- Todorov A, Engell AD (2008) The role of the amygdala in implicit evalua-

- tion of emotionally neutral faces. *Soc Cogn Affect Neurosci* 3:303–312. [CrossRef Medline](#)
- Todorov A, Pakrashi M, Oosterhof NN (2009) Evaluating faces on trustworthiness after minimal time exposure. *Soc Cogn* 27:813–833. [CrossRef](#)
- Todorov A, Said CP, Oosterhof NN, Engell AD (2011) Task-invariant brain responses to the social value of faces. *J Cogn Neurosci* 23:2766–2781. [CrossRef Medline](#)
- Todorov A, Mende-Siedlecki P, Dotsch R (2013) Social judgments from faces. *Curr Opin Neurobiol* 23:373–380. [CrossRef Medline](#)
- Whalen PJ (1998) Fear, vigilance, and ambiguity: initial neuroimaging studies of the human amygdala. *Curr Dir Psychol Sci* 7:177–188. [CrossRef](#)
- Whalen PJ, Rauch SL, Etcoff NL, McInerney SC, Lee MB, Jenike MA (1998) Masked presentations of emotional facial expressions modulate amygdala activity without explicit knowledge. *J Neurosci* 18:411–418. [Medline](#)
- Whalen PJ, Kagan J, Cook RG, Davis FC, Kim H, Polis S, McLaren DG, Somerville LH, McLean AA, Maxwell JS, Johnstone T (2004) Human amygdala responsivity to masked fearful eye whites. *Science* 306:2061. [CrossRef Medline](#)
- Williams LM, Das P, Liddell BJ, Kemp AH, Rennie CJ, Gordon E (2006) Mode of functional connectivity in amygdala pathways dissociates level of awareness for signals of fear. *J Neurosci* 26:9264–9271. [CrossRef Medline](#)
- Willis J, Todorov A (2006) First impressions: making up your mind after a 100-ms exposure to a face. *Psychol Sci* 17:592–598. [CrossRef Medline](#)
- Winston JS, Strange BA, O'Doherty J, Dolan RJ (2002) Automatic and intentional brain responses during evaluation of trustworthiness of faces. *Nat Neurosci* 5:277–283. [CrossRef Medline](#)
- Yarkoni T, Poldrack RA, Nichols TE, Van Essen DC, Wager TD (2011) Large-scale automated synthesis of human functional neuroimaging data. *Nat Methods* 8:665–670. [CrossRef Medline](#)
- Zald DH (2003) The human amygdala and the emotional evaluation of sensory stimuli. *Brain Res Rev* 41:88–123. [CrossRef Medline](#)
- Zebrowitz LA, Montepare JM (2008) Social psychological face perception: why appearance matters. *Soc Personal Psychol Compass* 2:1497. [CrossRef Medline](#)
- Zeger SL, Liang KY (1986) Longitudinal data analysis for discrete and continuous outcomes. *Biometrics* 42:121–130. [CrossRef Medline](#)