A memory advantage for untrustworthy faces

Nicholas O. Rule [a,*], Michael L. Slepian [b], Nalini Ambady [c]

[a] Psychology Department, University of Toronto, 100 St. George Street, Toronto, ON, Canada M5S 3G3
[b] Tufts University, United States
[c] Stanford University, United States

ARTICLE INFO

ABSTRACT

Inferences of others' social traits from their faces can influence how we think and behave towards them, but little is known about how perceptions of people's traits may affect downstream cognitions, such as memory. Here we explored the relationship between targets' perceived social traits and how well they were remembered following a single brief perception, focusing primarily on inferences of trustworthiness. In Study 1, participants encoded high-consensus trustworthy and untrustworthy faces, showing significantly better memory for the latter group. Study 2 compared memory for faces rated high and low on a series of traits (dominance, facial maturity, likeability, and trustworthiness), and found that untrustworthy and unlikeable faces were remembered best, with no differences for the other traits. Finally, Study 3 compared information about trustworthiness from facial appearance and from behavioral descriptions. Untrustworthy targets were remembered better than trustworthy targets both from behavior and faces, though the effects were significantly stronger for the latter. Faces perceived as untrustworthy therefore appear to be remembered better than faces perceived as trustworthy. Consistent with ecological theories of perception, cues to trustworthiness from facial appearance may thus guide who is remembered and who is forgotten at first impression.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

Walking along a busy street or through a crowd, many of the people that pass by slip through our thoughts relatively unnoticed—or, at least, unremembered (Simons & Levin, 1998). But some of the faces in the crowd stick with us. In the present work, we examined why, after brief encounters, some faces are better remembered than others. In doing so, we focused on the role of social traits, particularly trustworthiness. Perceivers quickly and consistently evaluate others' social traits at first perception (Ambady, Bernieri, & Richeson, 2000). The attribution of a variety of social traits to others based on photos of their faces shows agreement across perceivers after just 50 ms of viewing time (Bar, Neta, & Linz, 2006; Rule, Ambady, & Adams, 2009; Willis & Todorov, 2006). Given the efficient and persistent nature of these judgments at brief perception, it stands to reason that they might influence subsequent processing of the individuals with effects on memory. The question as to what physical qualities support face memory is one that has been explored in some depth. One of the strongest indicators of whether a face is remembered is its distinctiveness (Vokey & Read, 1992). People with atypical appearances are better remembered than those who appear more typical, likely because the former stand out as salient (Light, Kayra-Stuart, & Hollander, 1979).

To date, much of the research examining memory differences based on social traits has utilized methods that direct participants' attention to particular aspects of the targets. Typically, participants are asked to view a series of individuals and to evaluate aspects of their physical appearance (e.g., facial attractiveness; Barclay & Lalumiere, 2006), or to make specific inferences about their character

* Corresponding author.
E-mail address: rule@psych.utoronto.ca (N.O. Rule).

(e.g., likeability; Bell & Buchner, 2010). For example, in one study, participants were asked to rate a series of targets on a host of traits, including honesty, after which they chose among a number of professions (half of which were criminal occupations; e.g., armed robber) that they believed the target individuals might occupy (Yarmey, 1993); subsequently, the participants were given a memory test. This was not random design, as the researchers were interested in measuring whether trustworthy "good guys" and untrustworthy "bad guys" might be remembered differently; the participants' judgments foreshadowed this and were expected to influence memory.

Indeed, trustworthiness, a fundamental dimension of person judgment (Oostehrof & Todorov, 2008), has received considerable attention in the study of person memory (e.g., Oda, 1997). For instance, some studies have shown that untrustworthy individuals are remembered better than trustworthy individuals (e.g., Bayliss & Tipper, 2006; Mealey, Daood, & Krage, 1996; Mueller, Thompson, & Vogel, 1988; Oda, 1997), even among children (Kinzler & Shutts, 2008). This is often credited to the possibility of an evolved "cheater detection module" in human cognition and behavior (Cosmides, 1989). Other research (Barclay & Lalumiere, 2006; Mehl & Buchner, 2008) has documented some flaws in the methods used in several of these earlier studies and has highlighted the importance of equating the verbal information used to convey trustworthy and untrustworthy behaviors, as well as the importance of the context created by this verbal information (see below; e.g., Bell & Buchner, 2011). Moreover, the bulk of previous work has largely relied upon descriptions of people, rather than the natural signals of trustworthiness inherent in their facial appearances. Yet people show high agreement in their ratings of trustworthiness from faces (e.g., Zebrowitz, Voinescu, & Collins, 1996).

Asking participants to form impressions about trustworthiness and closely related traits (e.g., likeability; see Rule et al., 2010) prior to testing their memory may have limited ecological validity. In everyday life, one typically meets people without conscious goals to evaluate those individuals along a specific trait. Rather, our impressions of others are usually formed quickly, automatically, and without much conscious deliberation (Macrae & Bodenhausen, 2000). This is particularly relevant to cases of eyewitness identifications—an important area of research in memory (e.g., Meissner & Brigham, 2001). Witnessing a crime is generally fast, stressful, and rarely is one expected to be considering specific physical qualities or character traits of the perpetrator. Instead, our exposures to strangers—within and outside of crime scenes—tend to be brief, passing, and unconstrained by evaluative instructions. Thus, in the present work, we sought to explore memory for trustworthy versus untrustworthy targets using an experimental design that might be better-suited to understanding how impressions are formed in a real-world setting. We therefore exposed participants to faces in a passive-viewing task without any contextual information (such as behavioral descriptions; Mehl & Buchner, 2008), complex strategy evaluations (e.g., economic games; Oda & Nakajima, 2010), or instruction to assess specific traits related to the dimension upon which the targets were intended to differ (e.g., criminality, trustworthiness, likeability; Bell & Buchner, 2011); as used in previous studies.

The present work therefore sought to consider the role of facial trustworthiness in memory for individuals in a different context. Principally, we were interested in the role that facial appearance alone plays in memory for individuals seen once and without instruction to process the faces in any particular way; i.e., similar to encountering strangers in everyday life. In addition, we investigated the relative contributions of facial and verbal sources of information about trustworthiness to memory for faces. We therefore adopted an ecological approach to considering memory differences for glimpsed faces differing in trustworthiness.

Drawing from the ecological theory of object perception (Gibson, 1979), ecological theories of social perception posit that percepts have utility and function (McArthur & Baron, 1983; Zebrowitz & Collins, 1997). Facial displays are thought to have relevant inherent value for communicating properties about individuals, particularly their social traits (see Zebrowitz, 1997), and memory for faces might vary according to their functional value (see also Nairne & Pandeirada, 2008). Accordingly, trustworthy faces could signal opportunities for cooperation and affiliation, and untrustworthy faces might signal dangers to be avoided (e.g., Buchner, Bell, Mehl, & Musch, 2009; Cosmides, 1989; Slepian, Young, Rule, Weisbuch, & Ambady, 2012; Suzuki & Suga, 2010). From the perspective of the ecological theory of social perception, we expected untrustworthy faces to be salient to perceivers and therefore better remembered—both because they are considered to be of high ecological value (e.g., Buchner et al., 2009; Suzuki & Suga, 2010) and because they are likely to occur more rarely (e.g., Barclay, 2008), even when they are glimpsed briefly.

Throughout the evolution of research on the relationship between targets' trustworthiness and how well they are remembered, a number of nuances have emerged. An important theme seems to be that context is highly influential upon these effects. Initial accounts of differences in memory for individuals based on trustworthiness were situated within theories of an evolved capacity for "cheater detection" and reported differences in recognition memory between the trustworthy and untrustworthy targets (e.g., Mealey et al., 1996). Buchner and colleagues revised this perspective to suggest that source memory (the context within which a target is encountered) may be more valuable for evolutionary reasons than recognition memory, and presented empirical evidence in support of this claim (Bell & Buchner, 2010, 2011; Buchner et al., 2009).

Suzuki and Suga (2010) found that trustworthy faces were remembered better than untrustworthy faces but only when they were encoded in an incongruent context (i.e., as behaving in an untrustworthy manner), which participants learned through repeated experiences with the targets in an economic game. Thus, targets who essentially surprised participants with financial punishments were remembered best by the end of the game, suggesting that participants learned to become vigilant about these "wolves in sheep's clothing." Thus, the context within which targets were encoded into memory played an

important role in each of these effects, and depended critically upon the information that the experimenters supplied to the participants about the targets.

The current work therefore extended this consideration of context to a new domain with high ecological validity: memory for faces encountered briefly, as in first impressions of strangers. That is, similar to making judgments of people encountered in passing on the street, does the relative trustworthiness of those persons affect who is remembered and who is not? A central goal of this work was therefore to assess memory for faces in unconstrained situations where information about targets was captured from their appearances, rather than provided by verbal descriptors or direct interactive experience—such as that obtained in an economic game. Distinct from the majority of previous work in this research area, we therefore implemented an undirected (i.e., passive-viewing) incidental-encoding method of introducing perceivers to the targets and measured how properties of the faces' inferred traits might relate to whether they are remembered. Given the importance of contextual information in influencing memory, however, we were also concerned with the interaction between perceived (appearance-based) and described (verbally-stated) trustworthy behavior, which might help to illuminate some of the differences observed in previous studies (e.g., Barclay & Lalumiere, 2006; Kinzler & Shutts, 2008; Mealey et al., 1996; Yarmey, 1993). Thus, we sought to consider both sources of information, with particular attention paid to contributions from facial appearance, and examined the relative contributions of facial and verbal information.

Two studies have explored the relationship between the relative trustworthiness of targets' actual (rather than described) behavior and perceivers' memory for those faces. Yamagishi, Tanida, Mashima, Shimoma, and Kanazawa (2003) found that participants remembered the faces of people who defected in a prisoner's dilemma game better than the faces of people who cooperated. This was based only on their facial appearance, as the participants were not given information about who had defected and who had cooperated. Yet, unexpectedly, participants were unable to accurately distinguish the defectors from the cooperators when asked explicitly (as cited in Verplaetse, Vanneste, & Braeckman, 2007). Verplaetse et al. (2007) conducted similar studies in which they found that defectors and cooperators could be distinguished, however, and that their faces affected perceivers' subsequent recognition memory. The present work therefore aimed to investigate how the perception of social traits might influence individuals' memory for others' faces based on a single brief perception. Aside from trustworthiness, there has been some evidence that likeability, a similar and highly-correlated trait (e.g., Rule et al., 2010), affects the recognition of faces. Additional work has shown that other traits, such as dominance (Mueller & Mazur, 1996) and facial maturity (Berry & McArthur, 1985), are particularly important for the social perception of faces and that these traits are especially relevant to how salient a face is to a perceiver (Zebrowitz, 1997). In the current study, then, we examined the relationship between perceivers' memory for faces and their impressions on these four traits: trustworthiness, likeability, dominance, and facial maturity. Given the reviewed empirical and theoretical work, we predicted that untrustworthy faces would be better remembered than trustworthy ones. Moreover, consistent with the ecological perspective of perception, we hypothesized that facial information might provide a strong contribution to recognition memory, possibly greater than the influence of description-based information, as faces are more perceptually and ecologically salient (Zebrowitz, 1997).

In Study 1, we therefore examined whether faces perceived to be untrustworthy with high consensus would be remembered in a passive-viewing incidental-encoding recognition memory task better than faces consensually perceived to be trustworthy. In Study 2, we expanded this investigation by exploring the potential role that the other traits (likeability, dominance, and facial maturity) might play in both perceivers' memory for faces, as well as the memorability of individual faces when aggregated across perceivers. Finally, in Study 3, we investigated the interaction between facial cues to trustworthiness and explicit descriptions of the (un)trustworthy character of presented targets.

2. Study 1

Previous studies have shown mixed results for the relationship between recognition memory and targets' trustworthiness. Most of these studies manipulated target trustworthiness by varying whether the individuals were described as trustworthy or untrustworthy. Consequently, few studies (e.g., Suzuki & Suga, 2010; Verplaetse et al., 2007) have considered the influence of perceived trustworthiness due to differences in the targets' facial appearances. Thus, given the high consensus typically observed for judgments of face trustworthiness and the importance of the face in a number of social cognitive processes (including person memory; Zebrowitz, 1997), in Study 1 we tested perceivers' memory for faces that were consensually perceived as trustworthy and untrustworthy. Specifically, we presented participants with emotionally-neutral faces of men and women pre-rated to be high and low in expressing trustworthiness and untrustworthiness in a passive-viewing incidental-encoding recognition test and measured whether trustworthy and untrustworthy faces would be remembered differently. Due to potential differences in memory for faces according to their sex (e.g., Cross, Cross, & Daly, 1971; Wright & Sladden, 2003), memory for men's and women's faces were examined separately (Studies 1A and 1B, respectively).

2.1. Study 1A

2.1.1. Method

2.1.1.1. Stimuli. Forty grayscale images of male faces with neutral expressions, sized at 300×300 pixels, served as our experimental stimuli. The images were controlled for distance, angle, and luminosity. Targets were selected from a set of 59 male faces rated for trustworthiness in a previous study (Rule et al., 2009). From these, we selected the 20 faces with the highest ($M = 4.39$, $SE = .06$) and lowest

(M = 2.88, SE = .08) trustworthiness scores, as rated along a 7-point scale (1 = *Not at all trustworthy*, 7 = *Very trustworthy*) to compose our trustworthy and untrustworthy face sets, respectively.

2.1.1.2. Procedure. Undergraduates (N = 40, 48% female) completed an incidental-encoding memory task in exchange for partial course credit; the task consisted of several stages. In the first stage, participants were instructed that a series of faces would be presented on a computer screen and that they were to passively view each face. Participants saw 20 faces in random order for 3 s each. Each face was preceded by a 500-ms fixation cross and succeeded by 500 ms of rest. Unbeknownst to the participants, half (10) of the faces were previously rated as untrustworthy, and the other half (10) were rated as trustworthy. After participants viewed all of the faces, they were presented with a second set of instructions asking them to work on a word-search puzzle. The topic of the puzzle was unrelated to the experiment and was identical for each participant. Participants worked on the puzzle for 2 min, after which they were then presented with new instructions informing them that they would be presented with a series of faces, some of which they had seen in the first section of the experiment and some of which they had not. These faces consisted of all 20 trustworthy and untrustworthy targets from the first part of the experiment and 20 new faces, half of which were also previously rated as trustworthy and untrustworthy. Participants were asked to indicate via key press whether or not they had seen the face in the previous portion of the experiment. Although the task was self-paced, participants were instructed to work as quickly and accurately as possible. Critically, the assignment of faces as encoded or novel was counterbalanced.

2.1.2. Results and discussion

Data were analyzed using the signal detection statistic d' (Green & Swets, 1966). Faces that were correctly recognized from the encoding stage were treated as hits ($M_{\text{Perceived-Trustworthy}}$ = .63, SD = .23; $M_{\text{Perceived-Untrustworthy}}$ = .63, SD = .21), whereas faces incorrectly recognized were treated as false alarms ($M_{\text{Perceived-Trustworthy}}$ = .32, SD = .22; $M_{\text{Perceived-Untrustworthy}}$ = .22, SD = .24). This generated two d' scores for each participant: recognition for perceived-trustworthy targets and recognition for perceived-untrustworthy targets. These scores were then compared via a paired t-test that revealed that participants had significantly better memory for perceived-untrustworthy faces (M = 1.33, SD = 1.33), relative to perceived-trustworthy faces (M = .99, SD = 1.18): $t(39)$ = 2.38, p = .02, r = .36. The criterion threshold (c), a measure of response bias, did not significantly differ between the perceived-trustworthy (M = .14, SD = .60) and perceived-untrustworthy (M = .30, SD = .32) faces: $t(29)$ = 1.84, p = .07, r = .28. The effects did not vary across participant sex here or in any of the studies reported in the manuscript; we therefore do not discuss participant sex effects in the current work.

The faces of perceived-untrustworthy men were remembered significantly better than the faces of perceived-trustworthy men. Given that women are stereotyped to be more trustworthy than men, however (e.g., Eagly & Karau, 2002; Friedman & Zebrowitz, 1992), it is possible that differences in memory for faces high and low in perceived trustworthiness may vary with target sex. We therefore extended the work with men's faces in Study 1A to women's faces in Study 1B.

2.2. Study 1B

2.2.1. Method

Undergraduates (N = 30, 50% female) engaged in an incidental-encoding recognition memory test similar to Study 1A in exchange for partial course credit. The only change was that, rather than viewing men's faces, the stimuli consisted of 40 photographs of women's faces, also posing neutral expressions and sized at 300 × 300 pixels. The photos were selected from an in-house database of faces previously rated on trustworthiness. Half of the photos were relatively high in perceived trustworthiness (M = 4.56, SE = .05) and the remaining half were relatively low in perceived trustworthiness (M = 3.43, SE = .05), as rated along a scale anchored at 1 (*Not at all trustworthy*) and 7 (*Very trustworthy*).

2.2.2. Results and discussion

Data were analyzed using signal detection following the same definitions and procedures as in Study 1A (Hits: $M_{\text{Perceived-Trustworthy}}$ = .67, SD = .21; $M_{\text{Perceived-Untrustworthy}}$ = .80, SD = .14; False alarms: $M_{\text{Perceived-Trustworthy}}$ = .15, SD = .13; $M_{\text{Perceived-Untrustworthy}}$ = .15, SD = .11). As with the men's faces in Study 1A, perceived-untrustworthy women's faces (M = 2.05, SD = .67) were remembered significantly better than perceived-trustworthy women's faces (M = 1.58, SD = .77): $t(29)$ = 2.58, p = .02, r = .43. Moreover, measures of response bias (c) showed that participants had a significantly more stringent response threshold for perceived-trustworthy (M = .31, SD = .39) versus perceived-untrustworthy (M = .10, SD = .28) faces: $t(29)$ = 2.32, p = .03, r = .40. Interestingly, the direction of this effect is opposite that for the men's faces in Study 1A. Among male targets, there was a marginally significant tendency for participants to hold a more stringent response threshold for perceived-untrustworthy faces than for perceived-trustworthy faces. The reverse effect here may be related to an overall disposition toward considering women to be generally more trustworthy than men (e.g., Friedman & Zebrowitz, 1992), whereby perceivers are more conservative in their responses for targets whose sex and stereotypical trustworthiness are congruent. Nevertheless, perceivers' enhanced memory for perceived-untrustworthy over perceived-trustworthy faces appears to extend to faces of both sexes.

3. Study 2

In Study 1, we found that the faces of men and women perceived as untrustworthy were remembered significantly better than the faces of men and women perceived as trustworthy. The relative trustworthiness of the faces in these experiments was determined based on the mean

consensus ratings of external judges and the faces were pre-selected to differ as high and low in perceived trustworthiness. A question left from those studies, however, is whether perceptions of trustworthiness are critical to perceivers' memory for faces, overall, or if face memory might be distinguished along other variables as well.

We therefore explored this in Study 2 by asking participants to engage in a passive-viewing incidental-encoding memory task with a larger set of men's (Study 2A) and women's (Study 2B) faces whose perceived trustworthiness was not predefined. Subsequent to the recognition phase of the task, we asked participants to rate each face along a series of traits known to be important to social perception: dominance, facial maturity, likeability, and trustworthiness (Fiske, Cuddy, Glick, & Xu, 2002; Rule et al., 2010). We then measured the participants' memory for the faces rated as high and low on each of these traits following a tertile split of their mean ratings averaged across the participants in the sample. In addition, in Study 2B, we considered the relationship between how well each individual face was remembered and how these individual face-memory scores corresponded to the ratings given for the four traits.

3.1. Study 2A

3.1.1. Method

Undergraduates ($N = 43$, 40% female) engaged in an incidental-encoding memory task in exchange for partial course credit, as in Study 1. Faces of 60 men, none of which were used in Study 1A, were randomly selected from an in-house database of faces. Participants viewed 30 of the faces with instructions to view each face for its duration on the computer screen (3 s), with each face preceded by 500 ms of rest and a 500-ms presentation of a fixation cross to orient attention and to signal that the next face was forthcoming. Once participants viewed all 30 faces, they completed a word-search puzzle for 2 min, which served as a filler-task. Participants were then presented with the same 30 faces plus 30 previously unseen distracter faces (the assignment of targets and foils was counterbalanced across participants) in a random order with instruction to indicate whether they had seen each face in the previous section.

Subsequently, participants rated each of the 60 faces along several traits: dominance, facial maturity, likeability, and—critically—trustworthiness. Target faces were presented in random order, within four randomly-ordered blocks corresponding to each of the four traits. All ratings were made along 7-point scales anchored at 1 (*Not at all* X) and 7 (*Very* X); all inter-rater reliability Cronbach's α's > .87.

3.1.2. Results and discussion

Examining the neural correlates of trustworthiness judgments, Engell, Haxby, and Todorov (2007) found that the consensus of trustworthiness scores given to faces were a better predictor of participants' neural responses than were their own individual judgments, presumably due to measurement error in participants' idiosyncratic judgments. Following that work, ratings for each trait were averaged across participants to yield a mean score on the four traits for each target in the present analysis. These averages were then rank-ordered by value and a tertile split was used for each trait such that the 20 faces with the highest scores and the 20 faces with the lowest scores were distinguished for each individual trait. These tertiles were then applied to the recognition memory data to redefine the stimulus sets such that we calculated separate recognition memory scores for targets that were rated high on each trait and for targets that were rated low on each trait. Although the number of targets and foils that were rated high and low along each of these traits varied (see Table 1), the counterbalancing of targets and foils across participants prevented an imbalance in the data. Moreover, statistical tests of the frequency of targets and foils across the tertiles and traits showed a non-significant difference between the actual and ideal distributions: $\chi^2(9, N = 160) = 4.00$, $p = .91$. Hits and false-alarms were calculated separately for the high-ranking and low-ranking faces within each trait. From these, we calculated separate measures of recognition memory (d') for the two sets of faces within each trait. The high and low tertiles significantly differed for each of the four traits (see Table 2).

Participants remembered faces that they rated low in likeability and trustworthiness significantly better than faces that they rated high in likeability [$t(41) = 3.18$, $p = .003$, $r = .46$] and trustworthiness [$t(41) = 2.62$, $p = .01$, $r = .38$]. However, participants' memory for the faces rated high in dominance and facial maturity did not differ from those that were rated low in dominance [$t(41) = 1.68$, $p = .10$, $r = .25$] and facial maturity [$t(41) = .64$, $p = .52$, $r = .10$]. Response bias did not differ between the high- and low-rated faces along any of the traits: all t's < 1.16, all p's > .25; see Table 1 for descriptive statistics.

Consistent with the general findings of Study 1, then, participants remembered perceived-untrustworthy faces significantly better than perceived-trustworthy faces. Similar effects were found for targets' perceived likeability, a closely related trait (e.g., Rule et al., 2010). Specifically, participants remembered the faces of people they rated as unlikeable significantly better than the faces of people they believed to be likeable. In confluence, the findings for perceived likeability and perceived trustworthiness suggest that individuals express better recognition memory for faces that they evaluate negatively. Study 2B considers these effects for women's faces.

3.2. Study 2B

3.2.1. Method

Undergraduates ($N = 29$, 45% female) participated in exchange for partial course credit. Participants followed identical procedures to those of Study 2A with the exception that 60 women's faces were used in place of men's faces. As above, the photos were selected randomly from an in-house database and none of them were used in Study 1B. Participants showed acceptable levels of inter-rater reliability in their ratings of the faces along the four traits (all Cronbach's α's > .80).

Table 1
Descriptive statistics for the signal detection analyses in Study 2.

		Lower tertile							Upper tertile						
		Target:Foil ratio	H	FA	d'	SD	c	SD	Target:Foil ratio	H	FA	d'	SD	c	SD
Study 2A	Dominance	11:9	.60	.26	1.04	.74	.23	.47	7:13	.58	.31	.81	.69	.17	.53
	Facial maturity	10:10	.60	.29	.95	.78	.18	.45	11:9	.61	.27	1.05	.80	.19	.51
	Likeability	10:10	.63	.28	1.04	.70	.15	.44	8:12	.57	.35	.65	.75	.12	.52
	Trustworthiness	8:12	.61	.27	1.00	.81	.21	.42	11:9	.56	.35	.65	.73	.14	.55
Study 2B	Dominance	8:12	.66	.14	1.61	1.42	.35	.47	11:9	.70	.19	1.51	1.46	.18	.31
	Facial maturity	10:10	.72	.13	1.84	.72	.26	.36	7:13	.67	.15	1.57	.70	.27	.32
	Likeability	8:12	.70	.15	1.71	.75	.28	.26	8:12	.66	.26	1.21	.71	.14	.42
	Trustworthiness	11:9	.73	.14	1.84	.67	.22	.30	11:9	.65	.22	1.30	.65	.18	.45

Note: Designations of stimuli as targets and foils, and thus the Target:Foil ratios, were counterbalanced across participants.

Table 2
Means, variances, and significance tests for the upper and lower tertiles of targets rated along dominance, facial maturity, likeability, and trustworthiness in Study 2.

		Lower tertile		Upper tertile			
		M	SE	M	SE	t	r
Study 2A	Dominance	3.44	.07	5.12	.09	13.96***	.91
	Facial Maturity	3.80	.10	5.47	.07	14.01***	.92
	Likeability	3.30	.11	4.90	.08	12.16***	.89
	Trustworthiness	3.31	.09	4.69	.07	12.26***	.89
Study 2B	Dominance	3.35	.04	4.54	.08	12.94***	.90
	Facial Maturity	3.51	.05	5.23	.08	17.18***	.94
	Likeability	3.37	.09	5.01	.05	15.59***	.93
	Trustworthiness	3.12	.08	4.87	.06	18.15***	.95

*** $p < .001$.

3.2.2. Results and discussion

As in Study 2A, ratings for each trait were averaged across participants and the mean scores were rank-ordered to conduct a tertile split. The ratings for the upper and lower tertiles significantly differed for each of the four traits (see Table 2) and the ratio of targets to foils for each trait, which was counterbalanced across participants, did not differ from chance: $\chi^2(9, N = 160) = 4.00$, $p = .91$; see Table 1. We then calculated hits, false-alarms, and recognition memory (d') scores for each of the high- and low-rated faces within each trait. Consistent with the findings for men's faces in Study 2A, the faces of women rated low in trustworthiness were remembered significantly better than the faces of women rated high in trustworthiness: $t(27) = 2.92$, $p = .007$, $r = .49$. Similar to the findings for men's faces in Study 2A, women rated low in likeability were remembered better than women rated high in likeability; however, this difference was not significant when the Bonferroni correction ($\alpha = .0125$) was applied to control for multiple comparisons: $t(27) = 2.46$, $p = .02$, $r = .43$. Recognition memory for dominance [$t(27) = 1.09$, $p = .28$, $r = .21$] and facial maturity [$t(27) = 1.74$, $p = .09$, $r = .32$] did not significantly differ. Moreover, response bias did not differ among any of the four traits: all t's < 1.15, all p's > .26.

Both here and in Study 2A, we divided the data into tertiles based upon participants' ratings of the faces along various traits. This allowed us to determine how individuals' memory differed among the faces according to how they were rated on a variety of measures. An alternative approach, however, is to calculate a score for the memorability of each face and to then correlate these memorability scores with the trait ratings. This items-based analysis would allow us to utilize all of the faces, rather than removing the middle third. Given that the assignment of faces as targets and foils was counterbalanced across participants, we were therefore able to calculate the rate at which a particular face was accurately remembered (i.e., hits) and falsely misremembered (i.e., false-alarms) for each *face* across participants and to derive a corresponding memorability score using d' for each face. We then correlated these memorability scores with the mean ratings given to the faces along the four traits. To maintain sufficient power for the correlation analysis, we conducted these tests on the aggregated sample of male and female targets.

Results of the items-analysis mirrored those of the more traditional, participant-based analysis. Face memorability was unrelated to ratings of dominance [$r(118) = -.12$, $p = .20$] and facial maturity [$r(118) = -.07$, $p = .45$] but was significantly negatively correlated with perceptions of trustworthiness [$r(118) = -.25$, $p = .005$] and likeability [$r(118) = -.23$, $p = .01$]; see Fig. 1. These effects remained significant for both perceived trustworthiness [$r(117) = -.28$, $p = .002$] and likeability [$r(117) = -.28$, $p = .002$] when controlling for target sex in a partial correlation. Moreover, we asked separate raters to judge the attractiveness ($N = 14$; Cronbach's $\alpha = .81$) and distinctiveness ($N = 25$; Cronbach's $\alpha = .70$) of the faces along a scale ranging from 1 (*Not at all attractive/distinctive*) to 7 (*Very attractive/distinctive*), respectively. Attractiveness was not
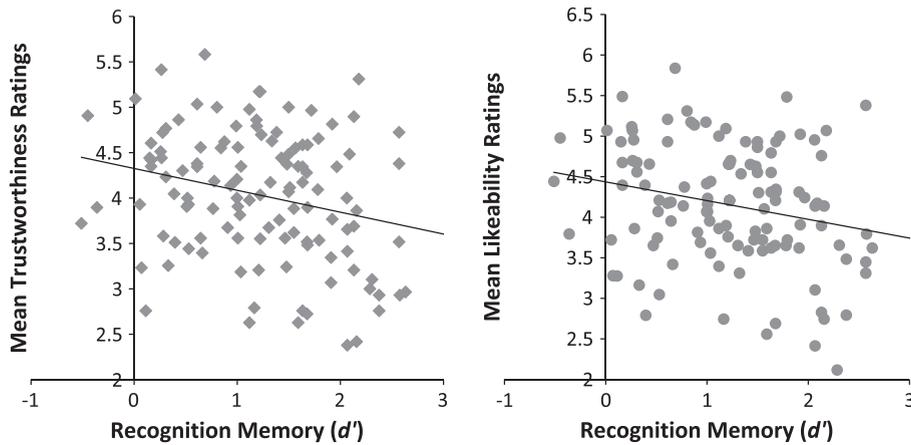
Fig. 1. Scattertplots depicting the relationship between recognition memory for individual target faces and mean consensus ratings of trustworthiness (left panel) and likeability (right panel) for the faces in Study 2.

related to either memory [$r(118) = .09$, $p = .33$] or response bias [$r(118) = -.14$, $p = .13$], whereas perceptions of distinctiveness were significantly correlated with both memory [$r(118) = .46$, $p < .001$] and response bias [$r(118) = .26$, $p = .004$], replicating past research (Valentine & Bruce, 1986; Vokey & Read, 1992). However, perceptions of likeability and trustworthiness were still significantly correlated with face memorability when attractiveness and distinctiveness were included as covariates in a partial correlation: r's < $-.24$, p's < $.01$. In addition, response bias was not correlated with any of the traits (all $|r|$'s < $.11$, all p's > $.22$), though the correlation with perceived likeability was only non-significant following Bonferroni-correction [$\alpha = .0125$; $r(118) = -.22$, $p = .02$], suggesting that a more stringent response criterion was applied to faces that were rated as less likeable. Indeed, although ratings of likeability and trustworthiness were significantly correlated [$r(118) = .80$, $p < .001$], as in previous work (e.g., Rule et al., 2010), they were not identical and ratings of trustworthiness were a slightly stronger predictor of face memory in the first-order correlation, as it was in the participant-based analysis for women's faces.

Negatively-valenced (i.e., perceived-untrustworthy, perceived-unlikeable) faces are therefore better remembered for both male and female targets, as predicted by the ecological account of face memory that we proposed (see also Bell & Buchner, 2010). When considering participants as the unit of analysis, individual judges remembered perceived as negatively-valenced better than faces perceived as positively-valenced. In turn, when considering targets as the unit of analysis, the more negatively-valenced that a face was perceived, the better it was remembered across the aggregated group of perceivers. In Study 3, we tested the interaction between appearance-based and description-based information about trustworthiness.

4. Study 3

Studies 1 and 2 showed that perceivers recognized faces perceived as untrustworthy with better accuracy

than faces perceived as trustworthy. In both sets of studies, however, information about trustworthiness was communicated through facial appearance. Although facial appearance is a remarkable resource for information about individuals' traits, behaviors, and social group memberships (e.g., Zebrowitz, 1997), facial appearance can also sometimes be misleading (Zebrowitz, Andreoletti, Collins, Lee, & Blumenthal, 1998). Moreover, we often acquire information about others' trustworthiness not just from how honest their faces appear but from their reputation and the observation of their behavior. Thus, external information about an individual's trustworthiness is a critical factor in whether we decide to trust that person.

In Study 3, then, we were interested in testing how participants' memory for perceived-trustworthy and perceived-untrustworthy faces might be affected by information about those persons. Specifically, the faces of perceived-trustworthy and perceived-untrustworthy men and women from Study 1 were paired with descriptive phrases that were either congruent or incongruent in trustworthiness. We then measured the participants' memory for the faces and examined the interaction between facial appearance and explicit information about trustworthiness. Based upon the ecological account proposed—wherein facial, relative to description-based, information is especially perceptually salient (Zebrowitz, 1997)—we hypothesized that individuals perceived as untrustworthy (versus trustworthy) from their faces would be better remembered than individuals described as untrustworthy (versus trustworthy) in their behaviors.

4.1. Method

Undergraduates ($N = 198$, 52% female) participated in exchange for partial course credit. Participants were randomly assigned to view either the men's ($n = 95$) or women's ($n = 103$) faces from Study 1. Of the 40 men's and 40 women's faces used in Studies 1A and 1B, respectively, half were high in perceived trustworthiness and the remaining half were low in perceived trustworthiness. In

the current study, we paired each of these faces with trustworthy and untrustworthy third-person descriptive sentences during the incidental-encoding stage of the experiment (see Appendix). The sentences were pre-tested for complexity and trustworthiness ($N = 20$; Cronbach's α's > .87). The two sentence types did not differ in complexity [$t(18) = .92$, $p = .37$] but sentences describing trustworthy behaviors were rated as significantly more trustworthy ($M = 2.36$, $SE = .10$) than those describing untrustworthy ($M = -2.25$, $SE = .13$) behaviors on a 7-point scale ranging from -3 (Untrustworthy) to 3 (Trustworthy): $t(18) = 27.05$, $p < .001$, $r = .99$. Comparisons of the absolute values of these ratings showed that the trustworthy and untrustworthy sentences did not significantly differ in their distance from the scale's midpoint (0): $t(18) = .80$, $p = .44$, $r = .18$. In addition, the groups of sentences did not significantly differ in terms of length, differing neither in the number of characters [$t(18) = .47$, $p = .64$] nor the number of words [$t(18) = .80$, $p = .43$].

Participants were randomly assigned to either encode the faces with congruent information (e.g., described-trustworthy sentences paired with perceived-trustworthy faces; $n = 51$ for female faces and $n = 45$ for male faces) or incongruent information (e.g., described-untrustworthy sentences paired with perceived-trustworthy faces; $n = 52$ for female faces and $n = 50$ for male faces). There were 10 trustworthy and 10 untrustworthy descriptive sentences, each of which was assigned to one face within condition (i.e., either congruent or incongruent) in a different randomized order for each participant. The stimuli and procedures were otherwise identical to Study 1. As above, the assignment of 20 faces as targets and 20 faces as foils was counterbalanced across participants. Thus, participants in each condition encoded 20 faces: 10 perceived-trustworthy faces described as trustworthy and 10 perceived-untrustworthy faces described as untrustworthy in the congruent condition, and 10 perceived-untrustworthy faces described as trustworthy and 10 perceived-trustworthy faces described as untrustworthy in the incongruent condition.

4.2. Results and discussion

Hits, false-alarms, discriminability (d'), and criterion (c) scores for recognition memory were calculated, as above, and tested via a 2 (target sex: male or female) × 2 (condition: congruent or incongruent faces and descriptors) × 2 (face trustworthiness: perceived-trustworthy or perceived-untrustworthy) ANOVA with repeated measures on the last factor. Similar to the effects found in Studies 1 and 2, perceived-untrustworthy faces were remembered significantly better than perceived-trustworthy faces: $F(1, 194) = 32.99$, $p < .001$, $r = .38$. This varied significantly by condition, however [$F(1,194) = 3.78$, $p = .05$, $r = .14$], whereby the memory advantage for perceived-untrustworthy faces was greater in the congruent [$t(95) = 5.45$, $p < .001$, $r = .49$] than incongruent [$t(101) = 2.67$, $p = .009$, $r = .26$] condition (Fig. 2). Thus, although the perceived trustworthiness of the face appeared to exert a large and consistent effect across both conditions, the information

accompanying the face also had a significant effect on whether the face was later recognized.

We also observed a main effect of target sex whereby women's faces ($M = 1.87$, $SE = .06$) were remembered significantly better than men's faces ($M = 1.67$, $SE = .06$): $F(1,194) = 5.85$, $p = .02$, $r = .17$. This was further qualified by a significant face trustworthiness × target sex interaction: $F(1,194) = 9.81$, $p = .002$, $r = .22$. As illustrated in Fig. 3, perceived-untrustworthy women were remembered better than each of the other three groups: all t's > 3.81, all p's < .001, all r's > .26. No other effects were significant: condition, the condition × target sex interaction, or the face trustworthiness × target sex × condition interaction: all F's < 1.47, all p's > .23.

Parallel analyses for participants' criterion scores showed only main effects for face trustworthiness [$F(1, 194) = 8.97$, $p = .003$, $r = .21$] and target sex: $F(1, 194) = 4.58$, $p = .03$, $r = .15$. The interaction between face trustworthiness and target sex did not reach significance [$F(1, 194) = 3.44$, $p = .07$] but the means reflected an effect like that observed in Study 1, wherein perceived-untrustworthy men and perceived-trustworthy women were more likely to be categorized with a stringent response criterion; all other effects in the model were nonsignificant: all F's < .97, all p's > .33. Participants held a significantly more stringent response criterion for perceived-untrustworthy ($M = .17$, $SD = .29$) versus perceived-trustworthy ($M = .08$, $SD = .36$) faces, and a significantly more stringent response criterion for male faces ($M = .17$, $SE = .02$) than female faces ($M = .09$, $SE = .03$).

Information about trustworthiness expressed by individuals' faces therefore exerts a strong, stable, and reliable influence upon perceivers' recognition memory. Explicit information about targets' trustworthiness is also important, however, as captured by the face trustworthiness × condition interaction. Although both the effects for face trustworthiness and sentence trustworthiness were statistically significant, the magnitude of these
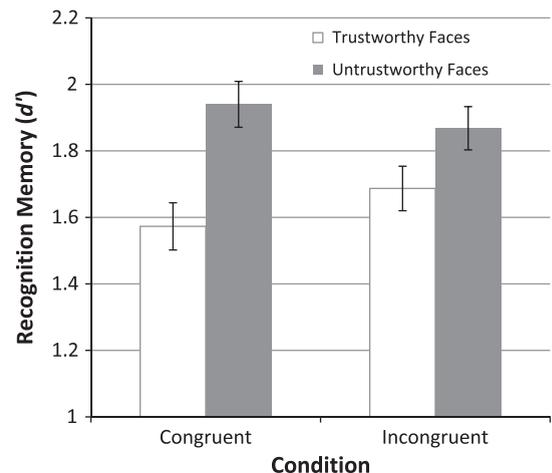
Fig. 2. Means and standard errors for participants' recognition memory of perceived-trustworthy faces and perceived-untrustworthy faces across the face-descriptor congruent and face-descriptor incongruent conditions in Study 3.
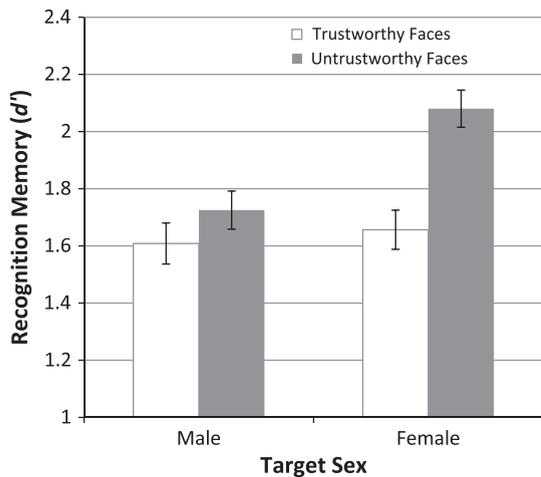
Fig. 3. Means and standard errors for participants' recognition memory of perceived-trustworthy and perceived-untrustworthy men's and women's faces in Study 3.

effects showed that the effect for face trustworthiness ($r = .38$) was significantly stronger than that for the information provided by the sentences ($r = .14$): $|Z| = 2.54$, $p = .006$.

5. General discussion

People appear to remember faces that are untrustworthy, or negatively-valenced, based on first impressions. In Study 1, participants remembered faces previously-rated as untrustworthy better than faces that had been previously rated as trustworthy. In Study 2, participants' memory for faces was strongly related to the trustworthiness and likeability of the faces, with better memory for perceived-untrustworthy and perceived-unlikeable faces. Finally, in Study 3, perceived-untrustworthy faces were remembered significantly better than perceived-trustworthy faces based on facial appearance alone. In addition, providing participants with information about target individuals' relative trustworthiness enhanced memory when the information was congruent with the face, but caused a reduction in memory when the apparent trustworthiness of the face did not match the trustworthiness of the information provided. Thus, trustworthiness appears to possess an important influence upon recognition memory for the faces of both male and female targets.

Most previous work examining memory and trustworthiness has manipulated targets' trustworthiness by providing participants with verbal information. These studies have shown some mixed effects, with some researchers reporting that individuals described as untrustworthy were remembered better than those described as trustworthy (e.g., Mealey et al., 1996), and others reporting null effects (e.g., Barclay & Lalumiere, 2006). Across three studies and two sets of male and female faces using a brief, passive-viewing memory task; here we found a strong and consistent effect by which facial appearance influences

memory. Although verbal information added to this effect, the independent contribution of the face was statistically greater. Thus, it is possible that information about trustworthiness from facial appearance might exert a stronger effect on recognition memory than does providing participants with verbal information, which tends to show more source memory effects (e.g., Buchner et al., 2009). Future work may wish to incorporate other aspects of the designs of previous studies (e.g., a longer retention interval, or considerations of source memory) for comparing facially- and verbally-based information on memory.

These data are also interesting in light of previous work that has shown a relationship between facial appearance and trustworthy behavior in an economic game. Specifically, Suzuki and Suga (2010) found that memory was best for the pairing of apparently trustworthy faces with untrustworthy behavior. In Study 3, however, we found that faces were better distinguished when they were accompanied by congruent information than when they were accompanied by incongruent information, and that this was true of both perceived-trustworthy and perceived-untrustworthy faces. One reason for the difference between the present results and those of Suzuki and Suga (2010) may be that the type of memory assessed was different. Here, we focused on incidentally-encoded recognition memory for faces that were seen once—a design meant to simulate brief encounters in the real world. Suzuki and Suga (2010), however, presented participants with the same faces over a series of trials during which the participants learned through experience whether particular targets could be trusted or not. It is thus reasonable that participants in that study would have come to remember trustworthy individuals who behaved in an untrustworthy manner because they represented occasions upon which participants would have needed to be especially vigilant, else they would lose money in the economic game. The learning component of Suzuki and Suga's (2010) work therefore distinguishes it from the "first impressions" nature of the current research. That said, women—who are generally believed to be more trustworthy than men (e.g., Friedman & Zebrowitz, 1992)—were remembered best when paired with untrustworthy information, similar to the "wolf in sheep's clothing" effect reported by Suzuki and Suga (2010).

One advantage of the present studies is that we observed convergent results when testing the relationship between face memory and trait information bidirectionally. That is, faces perceived as untrustworthy were remembered better than faces perceived as trustworthy both when the targets were pre-selected to differ in perceived trustworthiness a priori, as well as when face-memorability was related to targets' perceived traits post hoc. In Studies 1 and 3, faces consensually-regarded as untrustworthy were remembered better than faces consensually-regarded as trustworthy. In Study 2, splitting a group of faces based on high and low levels of various traits showed that trustworthiness and likeability (a highly correlated trait) uniquely predicted which faces were remembered and which were not. Moreover, the more likeable or trustworthy a particular face was, the better it was remembered by the participants as a group. Thus, recognition memory

was associated with trustworthiness both through experimenter and participant selection.

Another contribution of the current work is that we were able to distinguish the importance of the sources of information about trustworthiness upon participants' memory. Many of the previous studies have manipulated target trustworthiness using vignettes while controlling for contributions from facial appearance (e.g., Barclay & Lalumiere, 2006; Mealey et al., 1996). Here we observed that the face exerted a stronger influence on recognition memory than did experimentally-assigned information. Perceivers may therefore utilize information from scenarios less than information from facial appearance. Yet one important consideration is that our encoding task differed from that used in many of the previous studies. Whereas much of the earlier work asked participants to make judgments of the targets during encoding, here we instructed the participants to passively view the faces. Although the fact that we observed a significant effect even when the participants' attention was unconstrained by a task is impressive, it might possibly have caused them to rely more on facial information than on the behavioral descriptions. The significant effect of the descriptions suggests that participants did read and process the messages, yet the faces may have been more salient than the written information when presented without a directive task.

Recent work has posited a functional basis for memory (Nairne & Pandeirada, 2008). These accounts suggest that evolutionary pressures resulted in adaptations for memory processes in ways that assist survival. The present work adds to this literature by integrating data and theory from a functional perspective on perception that leads to differences in memory. Specifically, the ecological theory of perception suggests that individuals perceive objects and persons in terms of their functional value (e.g., Gibson, 1979; McArthur & Baron, 1983; Zebrowitz & Collins, 1997). Here, information about functionality was provided to participants in an indirect manner: principally through facial appearance but also through incidental-encoding of behavioral descriptions. These data could therefore be interpreted both within a functional account of memory, wherein untrustworthy targets are regarded as relevant for survival, and within a functional account of perception, wherein untrustworthy percepts are highly salient. Despite providing novel insights, the current work is not without limitations. One drawback is that we relied entirely on subjective assessments of trustworthiness. We do not know the true trustworthiness of any of the individual targets. Rather, we know only that some individuals are judged to be more trustworthy than others, that there is high consensus in these judgments, and that this belief has a powerful effect on whether the person is later remembered. Although some studies have linked face memory to targets' actual cheating behavior (Verplaetse et al., 2007; Yamagishi et al., 2003), cheating and trustworthiness are not entirely synonymous. Cheating may represent but one form of untrustworthy behavior, whereas trustworthiness may be a more general means of describing one's disposition or overall tendencies. Indeed, a difficulty inherent in such work is that boundaries for trustworthy versus untrustworthy behavior are

somewhat unclear. For example, do polite lies or sanctioned thefts render an individual as "untrustworthy," or must one's behaviors be more severe (e.g., stealing a large amount of money) or chronic (e.g., stealing pencils from work on a daily basis)? Without a strong operational definition of trustworthiness, it is difficult to test and reach conclusions about objective versus subjective assessments of trustworthiness and how they may, or may not, relate to other cognitive processes like recognition memory.

Similarly, little is known about the qualities that cause a face to appear trustworthy or untrustworthy. One model has suggested that perceptions of trustworthiness may be overgeneralized responses to emotional expressions (Oosterhof & Todorov, 2008). Specifically, computer modeling has shown structural overlap between faces perceived as trustworthy and expressions of happiness, and between faces perceived as untrustworthy and expressions of anger (see also Zebrowitz, 1997). This may be consistent with other research showing that facial appearance can come to reflect internal dispositions over time (Malatesta, Fiore, & Messina, 1987), and that men who behave honestly come to develop more honest appearances as they age (though the opposite is true for women; Zebrowitz et al., 1996).

In sum, the present work provides several novel methodological and theoretical contributions to the literature. Methodologically, we adopted a relatively conservative, passive-viewing incidental-encoding approach to measuring recognition memory. In so doing, we did not constrain or guide participants' perceptions of faces and relied principally on their own impressions of the faces absent explicit information about target trustworthiness. Under such a design, we still observed that targets consensually-perceived as untrustworthy were better remembered than those consensually-perceived as trustworthy. This allowed us to make several novel theoretical contributions as well. For example, in Study 3, we were able to compare contributions from facial appearance and the more common method of manipulating trustworthiness: providing participants with explicit verbal information. In this context, we found that facial appearance produced a larger effect than did descriptions of behavior. More important, this design allowed us to test and demonstrate the importance of information from the face on person memory in a way that fits with the common experience of briefly meeting strangers. This was guided by ecological theories of perception, which posit that perceptions have a functional basis. We may thus infer that (un)trustworthiness is a central factor governing whether we remember other people upon first encountering them, suggesting that this is an important feature in our assessment of others as it affects a basic cognitive process like memory. These data therefore underscore the prominent influence of the face in social and cognitive processing.

Acknowledgement

Appendix A

Trustworthy and untrustworthy behavioral descriptions used in Study 3.

Trustworthy descriptions	Untrustworthy descriptions
Is someone his (her) friends can always count on	Lies to his (her) parents
Always picks his (her) sister up from school on time	Takes advantage of his (her) friends
Is someone you can confide in	Took his (her) friend's book without asking
Never breaks a promise	Cheated on his (her) girlfriend (boyfriend)
Always goes out of his (her) way to help others	Stole money from his (her) mother's purse
Volunteers at the hospital after school	Hit a car in a parking lot and drove away
Helps his (her) friends who are in need	Doesn't keep his (her) friend's secrets
Cares about others	Betrayed his (her) best friend
Is always there for his (her) friends	Doesn't return things he (she) borrows
Is a reliable person	Is not a very honest person

References

Ambady, N., Bernieri, F. J., & Richeson, J. A. (2000). Toward a histology of social behavior: Judgmental accuracy from thin slices of the behavioral stream. *Advances in Experimental Social Psychology, 32*, 201–271.

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

Barclay, P. (2008). Enhanced recognition of defectors depends on their rarity. *Cognition, 107*, 817–828.

Barclay, P., & Lalumiere, M. L. (2006). Do people differentially remember cheaters? *Human Nature, 17*, 98–113.

Bayliss, A. P., & Tipper, S. P. (2006). Predictive gaze cues and personality judgments: Should eye trust you? *Psychological Science, 17*, 514–520.

Bell, R., & Buchner, A. (2010). Valence modulates source memory for faces. *Memory & Cognition, 38*, 29–41.

Bell, R., & Buchner, A. (2011). Source memory for faces is determined by their emotional evaluation. *Emotion, 11*, 249–261.

Berry, D. S., & McArthur, L. Z. (1985). Some components and consequences of a babyface. *Journal of Personality and Social Psychology, 48*, 312–323.

Buchner, A., Bell, R., Mehl, B., & Musch, J. (2009). No enhanced recognition memory, but better source memory for faces of cheaters. *Evolution and Human Behavior, 30*, 212–224.

Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with Wason Selection Task. *Cognition, 31*, 187–276.

Cross, J. F., Cross, J., & Daly, J. (1971). Sex, race, age, and beauty as factors in recognition of faces. *Perception & Psychophysics, 10*, 393–396.

Eagly, A. H., & Karau, S. J. (2002). Role congruity theory of prejudice toward female leaders. *Psychological Bulletin, 109*, 573–598.

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

Fiske, S. T., Cuddy, A. J., Glick, P., & Xu, J. (2002). A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology, 82*, 878–902.

Friedman, H., & Zebrowitz, L. A. (1992). The contribution of typical sex differences in facial maturity to sex role stereotypes. *Personality and Social Psychology Bulletin, 18*, 430–438.

Gibson, J. J. (1979). *The ecological approach to visual perception*. New York: Harper & Row.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.

Kinzler, K. D., & Shutts, K. (2008). Memory for "mean" over "nice": The influence of threat on children's face memory. *Cognition, 107*, 775–783.

Light, L. L., Kayra-Stuart, F., & Hollander, S. (1979). Recognition memory for typical and unusual faces. *Journal of Experimental Psychology: Human Learning and Memory, 5*, 212–228.

Macrae, C. N., & Bodenhausen, G. V. (2000). Social cognition: Thinking categorically about others. *Annual Review of Psychology, 51*, 93–120.

Malatesta, C. Z., Fiore, M. J., & Messina, J. J. (1987). Affect, personality, and facial expression characteristics of older people. *Psychology and Aging, 2*, 64–69.

McArthur, L. Z., & Baron, R. M. (1983). Toward an ecological theory of social perception. *Psychological Review, 90*, 215–238.

Mealey, L., Daood, C., & Krage, M. (1996). Enhanced memory for faces of cheaters. *Ethology and Sociobiology, 17*, 119–128.

Mehl, B., & Buchner, A. (2008). No enhanced memory for faces of cheaters. *Evolution and Human Behavior, 29*, 35–41.

Meissner, C. A., & Brigham, J. C. (2001). Thirty years of investigating the own-race bias in memory for faces: A meta-analytic review. *Psychology, Public Policy, & Law, 7*, 3–35.

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

Mueller, J. H., Thompson, W. B., & Vogel, J. M. (1988). Perceived honesty and face memory. *Personality and Social Psychology Bulletin, 14*, 114–124.

Nairne, J. S., & Pandeirada, J. N. S. (2008). Adaptive memory: Remembering with a stone-age brain. *Current Directions in Psychological Science, 17*, 239–243.

Oda, R. (1997). Biased face recognition in the prisoner's dilemma game. *Evolution and Human Behavior, 18*, 309–317.

Oda, R., & Nakajima, S. (2010). Biased face recognition in the Faith Game. *Evolution and Human Behavior, 31*, 118–122.

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

Rule, N. O., Ambady, N., & Adams, R. B. Jr., (2009). Personality in perspective: Judgmental consistency across orientations of the face. *Perception, 38*, 1688–1699.

Rule, N. O., Ambady, N., Adams, R. B., Jr., Ozono, H., Nakashima, S., Yoshikawa, S., et al. (2010). Polling the face: Prediction and consensus across cultures. *Journal of Personality and Social Psychology, 98*, 1–15.

Simons, D. J., & Levin, D. T. (1998). Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin & Review, 5*, 644–649.

Slepian, M. L., Young, S. G., Rule, N. O., Weisbuch, M., & Ambady, N. (2012). Embodied impression formation: Social judgments and motor cues to approach and avoidance. *Social Cognition, 30*, 232–240.

Suzuki, A., & Suga, S. (2010). Enhanced memory for the wolf in sheep's clothing: Facial trustworthiness modulates face-trait associative memory. *Cognition, 117*, 224–229.

Valentine, T., & Bruce, V. (1986). The effects of distinctiveness in recognising and classifying faces. *Perception, 15*, 525–535.

Verplaetse, J., Vanneste, S., & Braeckman, J. (2007). You can judge a book by its cover: The sequel. A kernel of truth in predictive cheating detection. *Evolution and Human Behavior, 28*, 260–271.

Vokey, J. R., & Read, J. D. (1992). Familiarity, memorability, and the effect of typicality on the recognition of faces. *Memory & Cognition, 20*, 291–302.

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

Wright, D. B., & Sladden, B. (2003). An own gender bias and the importance of hair in face recognition. *Acta Psychological, 114*, 101–114.

Yamagishi, T., Tanida, S., Mashima, R., Shimoma, E., & Kanazawa, S. (2003). You can judge a book by its cover: Evidence that cheaters may look different from cooperators. *Evolution and Human Behavior, 24*, 290–301.

Yarmey, A. D. (1993). Stereotypes and recognition memory for faces and voices of good guys and bad guys. *Applied Cognitive Psychology, 7*, 419–431.

Zebrowitz, L. A. (1997). *Reading faces: Window to the soul?* Boulder, CO: Westview Press.

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

Zebrowitz, L. A., & Collins, M. A. (1997). Accurate social perception at zero acquaintance: The affordances of a Gibsonian approach. *Personality and Social Psychology Review, 1*, 204–223.

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