Examining Interactive Effects of Group Membership and Untrustworthiness on Recognition Memory

Spring 2021 - Fall 2021

Faculty Mentor: Dr. Brittany Cassidy

Student: Grace Clark

Abstract

Faces provide certain cues, such as untrustworthiness, that have been shown to be more memorable than others. To better understand the strength of this untrustworthiness advantage in recognition memory, it is important to study faces that vary in trustworthiness and provide other cues known to affect recognition memory at the same time. This research measured recognition memory of faces that varied in group membership and trustworthiness. Participants took a bogus personality test to establish their relative ingroup and outgroup. Then, they completed an encoding phase in which they viewed ingroup and outgroup faces that were either trustworthy or untrustworthy. Because ingroup faces have been shown to be especially memorable, the untrustworthy advantage was expected to disappear for ingroup trustworthy faces. In addition, untrustworthy outgroup faces were expected to be remembered more than outgroup trustworthy and ingroup untrustworthy faces. These hypotheses were not supported, however. The current research builds on previous findings by investigating how two facial cues, trustworthiness and group membership, interact together to affect recognition memory.

Examining Interacting Effects of Group Membership and Trustworthiness on Recognition Memory

People are exposed to an incredible number of faces throughout their lifetimes and can recognize as many as 5000 of them (Jenkins, Dowsett, & Burton, 2018). Recognizing faces is important because it is key to facilitating social interactions and communication, as well as information about people's facial features. Some faces contain distinctive cues that have been empirically shown to stand out more than others. (Oosterhof & Todorev, 2008). One such cue is trustworthiness; a growing body of work shows that people remember untrustworthy faces more than trustworthy ones because they indicate potential danger (Rule, Slepian, & Ambady, 2012). Prior work examining facial untrustworthiness advantages in recognition memory has identified these advantages in the absence of other contextual cues. People do not always perceive faces in isolation from other cues potentially affecting recognition memory. People remember ingroup members' faces more than outgroup members' faces when they share group membership (Bernstein, Young, & Hugenberg, 2007).

To better understand untrustworthiness advantages in recognition memory, it is therefore important to investigate how people recognize faces varying in their trustworthiness when other contextual cues are available. The current work examined recognition advantages for untrustworthy faces identified as ingroup or outgroup members relative to perceivers. To this end, I first summarize literature on recognition advantages for untrustworthy faces. Then, I discuss literature on recognition advantages for ingroup faces. I then discuss the current research that examined how facial trustworthiness and group membership cues could interact to affect recognition memory.

Recognition Memory Advantages for Untrustworthy Faces

Recognizing untrustworthiness in others has important implications for survival because untrustworthiness informs people about whether it is dangerous to approach others (Hou & Liu, 2019). Furthermore, perception of untrustworthiness is believed to be so crucial for survival that according to evolutionary theory, people should remember survival related information better than other types of information (McBride, Thomas, & Zimmerman, 2013). Past recognition memory work supports this possibility (Rule et al., 2012). This prior work has sought to replicate the experience of quickly and passively viewing faces. The research consisted of three studies where people passively viewed faces rated as being highly trustworthy or highly untrustworthy. In Rule et al.'s (2012) first study, undergraduate participants completed an encoding task in which they were shown either male or female faces that had previously been evaluated as trustworthy or untrustworthy. Next, participants completed a distraction task. After a retention interval, participants completed a recognition task by identifying whether the previously seen trustworthy or untrustworthy faces were new.

Evidencing a recognition memory advantage for untrustworthy faces, participants had better recognition memory for untrustworthy versus trustworthy faces. These findings are consistent with other work that suggests that people are motivated to remember faces that give cues of danger to be avoided (Rule et al., 2012). The finding of a recognition advantage for untrustworthy faces provides a functional perspective within the ecological theory of perception; the ecological theory of perception posits that faces contain information that is crucial for communication about character traits (Rule et al., 2012). A functional perspective within this theory suggests that untrustworthy faces would be recognized more because this trait is considered more valuable for memory.

Recognition Memory Advantages for Ingroup Members

People remember more ingroup than outgroup faces. This finding has been found for racial ingroup and outgroup faces, and this is the basis of the cross-race effect (Young, Hugenberg, Bernstein, & Sacco, 2012). The cross-race effect refers to the tendency for people to remember own-race faces better than other-race faces (Young, et al., 2012). One influential theoretical explanation of the cross-race effect is the perceptual expertise model. This model suggests that people better recognize ingroup versus outgroup race faces because they process ingroup race faces more often (Young, et al., 2012). Social cognitive models elaborate on the perceptual expertise model by suggesting that perceptual expertise, social cognition, and motivation combine to explain own-race advantages in recognition memory. Social cognition theory explains that people think about outgroups in terms of their category, while they think about ingroup members as being more individual (Young, et al., 2012). In addition, motivation is thought to influence the amount of effort that people put into recognizing faces (Young, et al., 2012). Hybrid theories that combine elements of perceptual expertise, social cognition, and motivation are posited to be the best models to explain the cross-race effect (Young, et al., 2012). Hybrid theories provide a bigger framework for understanding the cross-race effect than the perceptual expertise or social cognitive theories alone.

As previously mentioned, the perceptual-expertise model has been the dominant explanation for the cross-race effect (Young, et al., 2012). However, other work has drawn on evidence of social categorization as a mechanism for causing own-group bias (Bernstein, Young, & Hugenberg, 2007). This work suggests that own-group bias is largely caused by thinking about the outgroup categorically, which is explained in the social cognitive model. In Bernstein et al.'s (2007) first study, participants saw faces belonging to their own university or a rival university. Simply assigning faces to a category elicited own-group bias. Unlike cross-race effect studies,

this categorization study held faces constant across the ingroup and outgroup. This study showed that categorization itself elicited own-group bias rather than a characteristic of the face. In the second study, they wanted to investigate whether non-pre-existing groups elicited better recognition memory, which was interpreted as own-group bias. Participants completed a bogus personality test before the learning phase of the experiment. The test gave the participants either a red or green personality, and they were also asked to wear a wristband that corresponded with their personality type. They found that group categorization brings about ingroup bias, and therefore found support for the social-categorization model of the cross-race effect. This finding is significant because it shows that non-pre-existing groups also bring about own-group bias.

Young, Bernstein, & Hugenberg (2010) examined the question of whether own-group biases occur during encoding and recognition. The Encoding Hypothesis motivated their predictions; this hypothesis suggests that own-group biases emerge when information is first processed in memory, which is called encoding. In Young et al.'s (2010) first study, they tested the Encoding Hypothesis by asking participants to complete a face recognition experiment in which they received either pre-encoding or post-encoding instructions asking them to pay special attention to the features of the faces. They found that participants given instructions prior to encoding did not show own-group bias, but those given instructions after this phase did show it. In the second study, they wanted to control when the in-groups and out-groups were created. The researchers found that an own-group bias only occurred when participants were given the personality type prior to encoding. Their findings support the hypothesis that own-group biases occur during encoding as opposed to post-encoding. Therefore, they found support for the Encoding Hypothesis by eliminating own-group bias in the first study and causing it in the second.

Current Research

Research has found that people remember untrustworthy faces more than trustworthy ones (Rule et al., 2012). In addition, researchers have found that when participants are assigned to made-up groups, there is an ingroup bias in recognition memory; they remember ingroup faces more than outgroup faces (Bernstein et al., 2007). To date, there is little research on how group membership and facial trustworthiness may interact to affect recognition memory. The current research investigated how well people remember trustworthy and untrustworthy faces that are arbitrarily categorized as belonging to an ingroup or outgroup. Consistent with prior work (Rule et al, 2012), this research used White male faces as stimuli in order to control for race and gender, which could otherwise have affected the results.

Prior work (Lazerus, Ingbretsen, Stolier, Freeman, & Cikara, 2016) suggests that group membership affects how emotions are interpreted. Research on the positivity bias has shown that people rate the emotional expressions of ingroup members more positively than they do for outgroup members (Lazerus et al., 2016). This study explored how assigning groups to people affects their perceptions of ingroup and outgroup members' emotional expressions. In Lazerus et al.'s (2016) first study, participants were assigned to teams to specifically test how groups affect judgments of emotions. Simply labeling an emotional face with ingroup membership resulted in a more positive perception. Thus, it seems plausible that an untrustworthy ingroup face could be perceived as less untrustworthy than an untrustworthy outgroup. This differential perception could affect the untrustworthy advantage in recognition memory.

Hypotheses

Consistent with identified recognition memory advantages for untrustworthy faces (Rule et al., 2012), I hypothesized that people would better recognize untrustworthy versus trustworthy

faces. Consistent with identified recognition memory advantages for same-category faces (Bernstein et al., 2007), I hypothesized that people would better remember faces categorized as ingroup versus outgroup members. Critically, I also hypothesized that facial trustworthiness and group membership would interact to affect recognition memory. Specifically, I hypothesized that because ingroup untrustworthy faces are likely to be more positively perceived (Lazerus et al., 2016), that untrustworthiness advantages in recognition memory would be stronger for outgroup versus ingroup faces.

An alternative hypothesis was that untrustworthy faces would be so salient to perceivers (Mealey, Daood, & Krage, 1996) that ingroup membership would not reduce the untrustworthy advantage in recognition memory. Rather, outgroup membership could exacerbate this recognition advantage. This alternative pattern was motivated by work showing that people recognize the faces of those who cheat more than trustworthy faces because the former could signal danger or harm (Mealey, et al., 1996). An additional possibility was that the ingroup untrustworthy faces and the outgroup trustworthy faces would violate the expectations of participants. This was motivated by past work (Suzuki & Saga, 2010) showing that people's memory is enhanced when they encounter trustworthy-looking cheaters. However, the current research measured recognition memory, while Suzuki & Saga (2010) measured a different form of memory in their research.

Method

Participants

One hundred and forty-nine adult participants were recruited from the Amazon Mechanical Turk and given $1 for their participation. When taking exclusion criteria into account, the participant number decreased significantly. Participant number decreased to 143

when accounting for the manipulation checks, which consisted of asking participants their personality type before the encoding task as well as before the demographics section. In addition, participants who answered less than 7 distractor task questions were removed because of concerns that the stimuli would not transfer to long-term memory for the retrieval task. This brought the participant number down to 135. Participants whose accuracy on the distractor task was less than 75% were also removed for concerns about memory storage, and this brought the total number down to 129. When participants who did not respond to at least 85% of encoding task questions were removed, participant number dropped to 125. Participants who were less than 75% accurate on the encoding task were also removed, bringing the total number to 91. These two exclusion criteria were implemented because sufficient encoding is necessary for information to be stored in memory. If participants did not see most of the stimuli and accurately identify their personality type, neither trustworthiness nor group membership would be manipulated. Finally, participants who had less than a 50% hit rate (i.e., chance level memory performance) on the recognition task were removed, yielding a final analyzed sample of 58 participants (M_{age}=38.93 years, SD=10.79; $M_{years\ of\ education}$ = 14.97 years, SD = 2.14). Fifty-five participants identified as White, two participants identified as Asian, and one identified as Black. Of the 58 participants, 55 also identified as non-Hispanic.

Materials

Faces. Eighty-eight neutrally expressive White male faces were selected from the Chicago Face Database (Ma, Correll, & Wittenbrink, 2015). Trustworthiness norms from the database were used to classify faces as trustworthy or untrustworthy. Half of the faces were randomly assigned to have a red or blue personality type. These faces were equally distributed among faces assigned to be target (“old”) or lure (“new”) faces. Four task versions

counterbalanced personality type and task assignment of the faces. A 2 (Trustworthiness: trustworthy, untrustworthy) x 2 (Task assignment: target, lure) x 2 (Personality: red, blue) ANOVA on facial trustworthiness norms only elicited a main effect of trustworthiness, $F(1, 80) = 172.95$, $p < .001$, showing trustworthy faces were more trustworthy than untrustworthy faces. Selected faces were grey-scaled and cropped to the face area (e.g., hair and clothing cues were removed).

Personality Test. Participants rated their agreement (1 = strongly disagree, 7 = strongly agree) with questions from the Ten-Item Personality Inventory (Gosling, Rentfrow, & Swann, 2003).

Design

After taking the personality test, participants were randomly assigned (unbeknownst to them) to having a red or blue personality type. Untrustworthy and trustworthy faces were also randomly assigned as being labeled as having a red or blue personality type. By randomly assigning participants to one of two personality types, this created their ingroup and outgroup. Participants were exposed to trustworthy and untrustworthy ingroup and outgroup faces, and the proposed research was concerned with the group membership of the faces rather than the personality type assigned to participants. Participants saw faces across the four combinations of facial trustworthiness and group membership (trustworthy/ingroup, trustworthy/outgroup, untrustworthy/ingroup, untrustworthy outgroup). Thus, the primary analyses examining recognition memory advantages employed a 2 (Target Group Membership: ingroup, outgroup) x 2 (Target Facial Trustworthiness: trustworthy, untrustworthy) within-subjects design.

Procedure

After providing informed consent, participants began the experiment. The experiment was conducted remotely on participants' computers. Participants took the above-described personality test. After taking the test, participants were randomly told that their results indicated that they had a red or blue personality type. Participants were not informed that their personality type was randomly determined. After receiving their results, participants answered a manipulation check for personality type. Participants then completed the learning phase of the experiment. Participants saw 44 faces one at a time in a random order. Faces were presented for 2 seconds each, with 1 second between each face. Participants responded with a button press to indicate the group type of the faces on each trial. Participants pressed 1 if the face was a red personality and 2 if it was a blue personality. The personality type was listed above each face, and each personality type was also listed below each face in the corresponding color with the appropriate number press (i.e., red personality (1), blue personality (2)). Of the 44 randomly selected faces, there were 11 faces across the four combinations of facial trustworthiness and group membership (trustworthy/ingroup, trustworthy/outgroup, untrustworthy/ingroup, untrustworthy outgroup).

After completing the learning phase, participants completed a distractor task in which they solved 8 simple math problems with 8 seconds allotted for each question. After the distractor task, participants completed the recognition phase of the experiment. In the recognition phase, participants saw 88 faces. Forty-four of the faces were seen during the learning phase. The other 44 faces were new faces that participants did not see during the learning phase. The new 44 faces were evenly distributed across trustworthiness/untrustworthiness and red/blue personality type. The recognition phase was self-paced. Participants pressed 1 to indicate that they have seen the face before (i.e., it was an "old" face) and pressed 2 to indicate that they had not seen the face

before (i.e., it was a new face). There were eight task versions counterbalancing whether trustworthy and untrustworthy faces were assigned to red or blue personality types and whether they were used in the learning phase. After completing the recognition phase, participants answered another manipulation check for personality type as well as a brief demographics section asking for age, education level, race, and ethnicity. After finishing the experiment, participants were thanked and debriefed.

Results

D'

Like prior work (Bernstein et al., 2007), recognition memory was measured using d'. D' is a measure of sensitivity, or the distance between the means of the hit and false alarm rate. Previously seen ("old") faces identified as "old" are hits. New faces incorrectly identified as "old" are false alarms. To find d', the standardized difference between the hit rate and false alarm rate was calculated (Stanislaw & Todorov, 1999). Larger values for d' indicate greater sensitivity, or a greater difference between the hit and false alarm rate. A d' value close to zero indicates that participants performed at chance. D' values were calculated separately for untrustworthy and trustworthy ingroup and outgroup faces. "Old" untrustworthy ingroup faces identified as such were considered hits, while new untrustworthy ingroup faces identified as "old" were considered false alarms. D' was used to analyze the results rather than relying only on hit rate because it corrects for response bias, or the tendency for participants to randomly answer questions (Stanislaw & Todorov, 1999).

A two-way repeated measures ANOVA for D' found a marginal Facial Trustworthiness effect, $F(1, 57) = 2.696$, $p = 0.106$, $\eta_p^2 = 0.007$. Consistent with past research (Rule et al., 2012), people had more sensitivity in recognizing untrustworthy relative to trustworthy faces. There was

no Group Membership effect, $F(1, 57) = 0.178$, $p = 0.675$, $\eta_p^2 = 0.0004$, or Facial Trustworthiness x Group Membership interaction, $F(1, 57) = 0.522$, $p = 0.473$, $\eta_p^2 = 0.001$. See Table 1 for descriptive statistics on Group Membership and Trustworthy categories for D'.

Hits

As described above, hits refer to previously seen ("old") faces that are correctly identified as "old." A two-way repeated measures ANOVA for Hit Rate found no significant effect for Facial Trustworthiness, $F(1, 57) = 0.138$, $p = 0.712$, $\eta_p^2 = 2.967 \times 10^{-4}$, no effect for Group Membership, $F(1, 57) = 0.776$, $p = 0.382$, $\eta_p^2 = 3.972 \times 10^{-3}$, and no Trustworthiness x Group Membership interaction, $F(1, 57) = 0.0005$, $p = 0.983$, $\eta_p^2 = 1.299 \times 10^{-6}$. See Table 2 for descriptive statistics for Trustworthiness and Group Membership categories for Hit Rate.

False Alarms

As described above, false alarms refer to "new" faces that are incorrectly identified as "old" faces. A two-way repeated measures ANOVA for False Alarm Rate found a marginal effect for Facial Trustworthiness, where trustworthiness was remembered more than untrustworthiness, $F(1, 57) = 2.376$, $p = 0.129$, $\eta_p^2 = 0.007$. There was no Group Membership effect, $F(1, 57) = 0.288$, $p = 0.594$, $\eta_p^2 = 0.001$, or Facial Trustworthiness x Group Membership interaction, $F(1, 57) = 0.327$, $p = 0.590$, $\eta_p^2 = 0.0008$. See Table 3 for descriptive statistics for Trustworthiness and Group Membership categories for False Alarm Rate.

Discussion

It was predicted that untrustworthy faces would be remembered more than trustworthy faces and that ingroup faces would be remembered more than outgroup faces. Additionally, I hypothesized that facial trustworthiness and group membership would interact to affect group

membership, and that the untrustworthiness advantage would be reduced for ingroup faces. Overall, the results are not in line with the hypotheses. The results show that the untrustworthiness effect was likely maintained across group contexts; this effect is marginal, but it is in line with my prediction as well as past research (Rule et al., 2012). The results did not find a general group membership effect, which is not in line with the hypotheses or past research findings that ingroups are remembered more than outgroups (Bernstein et al, 2017). This may be because trustworthiness is a more powerful cue than group membership. Another possibility is that the group membership cue used in the current research, the Ten-Item Personality Inventory (Gosling, Rentfrow, & Swann, 2003), did not elicit sufficient group manipulation. A potential mitigation for the lack of group membership salience is to ask participants to write down their personality type as well as verbally communicate it to the researcher. These steps would presumably increase the salience of the bogus personality test. In addition, if replicating this study online, it would be useful to use a group membership cue that is known to elicit strong ingroup and outgroup bias, rather than an arbitrary personality test.

It was predicted that because ingroups are rated more positively than outgroups, this would cause a diminished untrustworthy advantage among ingroup faces. However, no interaction was found for a reduced untrustworthy advantage in the ingroup, indicating that group membership does not have a significant influence on trustworthiness. Untrustworthiness cues may be more salient than group category cues when they are presented together, and this possibility counters the hypotheses of the current research. Speculatively, a potential explanation for the lack of interaction between facial trustworthiness and group membership could be explained by evolutionary theory, which posits that trustworthiness perception is crucial to survival (McBride, Thomas, & Zimmerman, 2013). It could be that when presented with

trustworthy and group membership cues, people give more sway to the trustworthy cues because they are deemed more relevant for survival, and therefore convey more important information than group membership cues. Future work should examine this possibility.

A significant limitation to the current research was the participant exclusion, which limited the statistical power of the study. The participant exclusion removed participants who did not pass the manipulation check, did not answer at least 85% of encoding questions, and did not have at least 75% accuracy on the retrieval task. This exclusion brought the participant number down from 149 to 58. This study was conducted online rather than in a lab, and the online format could be a contributor to the high number of participants who were excluded. The format of this online study could have resulted in a greater likelihood for distraction than would be expected in a study conducted in a lab. Replicating the current research in a lab rather than online would be useful in order to investigate differences in the distraction of participants. In other words, it seems likely that participants would be less distracted if they completed the study in the lab, and thus accuracy in the retrieval task would increase.

Another limitation to the current research is that the stimuli were White male faces. This was necessary in order to control for other variables such as race and gender that might otherwise impact the findings. The current research sought to measure how recognition memory is affected only by group membership and trustworthiness, not race or gender. Thus, it is unclear how the findings generalize across female faces and to minority faces. Investigating how these cues affect minority faces is also important because racial stereotypes could be linked to untrustworthiness, which could in turn lead to a greater untrustworthy advantage for the outgroup. The current research used arbitrary group membership; however, using group membership cues such as race or political affiliation could lead to more salient group membership bias. Real-life group

memberships like race and political affiliation are associated with stereotypes, while the arbitrary group membership used in the current research is not associated with existing stereotypes. Also, future research should investigate how group membership and facial characteristics affect recognition memory of women's faces because women are perceived differently in terms of their social traits (Vianello, Schnabel, Sriram, & Nosek, 2013).

Additionally, future research could investigate how other facial characteristics such as likeability and dominance as well as group membership affect recognition memory. This is a valuable avenue for research because recognizing these characteristics is important during social interactions (Oosterhof & Todorev, 2008). Another direction for future research is to test the recognition memory of participants without color context at recognition. This is an important next step to take because it is valuable to know whether participants remember the faces because of the color context or because of specific facial features. Overall, the current research takes an important step in investigating how more than one characteristic affects recognition memory.

References

Bernstein, M. J., Young, S. G., & Hugenberg, K. (2007). The cross-category effect: Mere social categorization is sufficient to elicit an own-group bias in face recognition. *Psychological Science*, *18*(8), 706-712.

Brewer, M. B. (1999). The Psychology of Prejudice: Ingroup Love or Outgroup Hate? *Journal of Social Issues, 55*(3), 429-444.

Cesario, J., Plaks, J. E., & Higgins, E. T. (2006). Automatic social behavior as motivated preparation to interact. *Journal of Personality and Social Psychology, 90*(6), 893–910.

Hou, C., & Liu, Z. (2019). The Survival Processing Advantage of Face: The Memorization of the (Un)Trustworthy Face Contributes More to Survival Adaptation. *Evolutionary Psychology, 17*(2), 1-12.

Hutchings, R. J., Simpson, A. J., Sherman, J. W., & Todd A. R. (2021). Perspective taking reduces intergroup bias in visual representations of faces. *Cognition*, *214*, 0010-0277.

Jenkins, R., Dowsett, A. J., & Burton, A. M. (2018). How many faces do people know?. *Proceedings. Biological sciences*, *285*(1888), 20181319.

Lazerus, T., Ingbretsen, Z. A., Stolier, R. M., Freeman, J. B., & Cikara, M. (2016, October 24). Positivity Bias in Judging Ingroup Members' Emotional Expressions. Emotion. Advance online publication.

Ma, D.S., Correll, J. & Wittenbrink, B. (2015). The Chicago face database: A free stimulus set of faces and norming data. *Behav Res* 47, 1122–1135.

McBride, D. M., Thomas, B. J., & Zimmerman, C. (2013). A test of the survival processing advantage in implicit and explicit memory tests. *Memory & Cognition, 41*(6), 862–871.

Mealey, L., Daood, C., & Krage, M. (1996). Enhanced Memory for Faces of Cheaters. *Ethology*

and Sociobiology, 17, 119-128.

Oosterhof, N. N., & Todorev, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences of the United States of America, 105*(32), 11087-11092.

Rule, N. O., Slepian, M. L., & Ambady, N. (2012). A memory advantage for untrustworthy faces. *Cognition*, *125*(2), 207-218.

Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments, & Computers, 31*(1), 137-149.

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.

Vilanello, M., Schnabel, K., Sriram, N., & Nosek, B. (2013). Gender differences in implicit and explicit personality traits. *Personality and Individual Differences, 55*(8), 994-999.

Young, S. G., Bernstein, M. J., & Hugenberg, K. (2010). When do own-group biases in face recognition occur? Encoding versus post-encoding. *Social Cognition*, *28*(2), 240-250.

Young, S. G., Hugenberg, K., Bernstein, M. J., & Sacco, D. F. (2012). Perception and motivation in face recognition: A critical review of theories of the cross-race effect. *Personality and Social Psychology Review*, *16*(2), 116-142.

Table 1

Means and Standard Deviations of D' ANOVAs for Trustworthiness and Group Membership Categories

Facial Trustworthiness	Group Membership	Mean	Standard Deviation
Trustworthy	Ingroup	0.790	0.710
Trustworthy	Outgroup	0.767	0.712
Untrustworthy	Ingroup	0.857	0.789
Untrustworthy	Outgroup	0.942	0.798

Table 2

Means and Standard Deviations of Hit ANOVAs for Trustworthiness and Group Membership Categories

Facial Trustworthiness	Group Membership	Mean	Standard Deviation
Trustworthy	Ingroup	0.639	0.171
Trustworthy	Outgroup	0.661	0.052
Untrustworthy	Ingroup	0.645	0.052
Untrustworthy	Outgroup	0.667	0.052

Table 3

Means and Standard Deviations of False Alarm ANOVAs for Trustworthiness and Group Membership Categories

Facial Trustworthiness	Group Membership	Mean	Standard Deviation
Trustworthy	Ingroup	0.374	0.184
Trustworthy	Outgroup	0.397	0.192
Untrustworthy	Ingroup	0.352	0.188
Untrustworthy	Outgroup	0.354	0.206