5-2022

Does Target Race Moderate the Effect of Trustworthiness on Face Memory?

Sharon Ada Kobinah

Abstract

Multiple factors can influence how accurately people remember the faces of others. Social group membership, and particularly race, is a well-studied influence. People tend to show better memory for racial ingroup faces than racial outgroup faces, a phenomenon known as the Cross Race Effect. Some other work has examined the influence of factors like facial trustworthiness in memory, finding that people may remember untrustworthy faces more accurately than trustworthy faces. The present study examines the joint influence of these two factors in order to determine whether target race moderates this untrustworthiness advantage in memory. White and Black participants encoded an equal amount of trustworthy and untrustworthy White and Black male faces. Although both participant groups accurately remembered ingroup untrustworthy faces better than trustworthy ones, only Black participants showed an untrustworthiness advantage for outgroup faces. Among White participants, this untrustworthiness advantage was limited to ingroup faces. These findings have implications for existing theory on appearance-based influences on face memory.

Keywords: cross race effect, untrustworthy advantage, facial memory, signal detection

MONTCLAIR STATE UNIVERSITY

Does Target Race Moderate the Effect of Trustworthiness on Face Memory?

by

Sharon Ada Kobinah

A Master's Thesis Submitted to the Faculty of

Montclair State University

In Partial Fulfillment of the Requirements

For the Degree of

Master of Arts

May 2022

College of Humanities and Social Sciences

Department of Psychology

Thesis Committee:

[redacted]

Dr. John Paul Wilson

Thesis Sponsor

[redacted]

Dr. Michael Bixter

Committee Member

[redacted]

Dr. Jennifer Yang

Committee Member

DOES TARGET RACE MODERATE THE EFFECT OF TRUSTWORTHINESS ON FACE MEMORY?

A THESIS

Submitted in partial fulfillment of the requirements

For the degree of Master of Arts

by

SHARON ADA KOBINAH

Montclair State University

Montclair, NJ

2022

Acknowledgments

I would like to extend a special thank you to my advisor, Dr. John Paul Wilson, for all of the years of support and encouragement. Thank you for working with my abnormal work schedule and going above and beyond to give me guidance every step of the way. I have a greater appreciation for this field than when we first began and I am thankful for everything I have learned throughout this study. Additionally, I would also like to thank my committee members, Dr. Michael Bixter and Dr. Jennifer Yang, for all of the insight, feedback, and support.

I would also like to thank my family and friends for their encouragement. An extended appreciation goes out to my parents. Words cannot express how grateful I am for your endless support in my endeavors. I am incredibly proud to have achieved this milestone in my life and I owe it to your sacrifices. To be the first to achieve this level of education in our family was made possible through your unwavering belief in me. Thank you for all that you do.

Table of Contents

List of Figures

Introduction

As humans, we are constantly attending to and remembering others in our social environment, and we do so more efficiently and accurately for some people than other people. There are certain factors that influence who we remember, and much research has focused on studying these factors. For example, our memory for others is impacted by their personality traits (Hastie & Kumar, 1979), the direction of their eye gaze (Mason et al., 2004), their race (Meissner & Brigham, 2001), ingroup/outgroup distinctions (Bernstein et al., 2007), and even factors as basic as facial appearance (Rule et al., 2012).

There are instances where we are incorrect in our facial recognition, despite the steps and cues we take to try and recall. Whether trying to remember how we recognize someone we've seen in a social setting or attempting to choose a suspect in a police line-up, there are occasions where our facial recall is not accurate as we would hope. In several circumstances, these misidentifications could have social and even serious consequences (Brigham & Malpass, 1985). This project will examine the intersection of two such factors: race and facial trustworthiness.

Cross-Race Effect

People tend to be more accurate at remembering the faces of members of their own race than members of other races (Bothwell et al., 1989; Chiroro & Valentine, 1995; Tanaka et al., 2004). This cross-race effect (CRE) in memory has been demonstrated in and out of the lab, and it is one of the more easily replicable findings in psychology. Recent theoretical perspectives on the CRE have suggested that it may not be race *per se* that creates this bias (Hugenberg et al., 2013). Rather, people tend to show

enhanced attention and memory for members of *ingroups*. That is, people better remember others who belong to their own social groups than people who belong to opposing social groups. Regardless of the exact cause of the CRE, though, race is very clearly one dimension upon which memory biases occur.

There have been decades of research demonstrating that people are better at recognizing the faces of ingroup members than outgroup members. Meta-analyses have demonstrated that this is a robust phenomenon (Meissner & Brigham, 2001) that occurs across many different participants and target racial groups and among people of varying ages. However, it is worth noting that some studies show that the CRE may not occur to the same magnitude for all groups. For example, Feinman & Entwisle (1976) found that young Black participants did not show as strong of a CRE as young White participants. Meta-analytically, there is also some evidence that majority group members may show a larger CRE than minority group members (Meissner & Brigham, 2001).

Although some perspectives have argued that the CRE occurs due to a lack of "perceptual expertise" with the faces of other-race people (e.g., Rhodes et al., 1989), others argue that the CRE is merely one manifestation of a broader tendency to misremember the faces of outgroup members. It can be argued then that categorical thinking within a social setting on outgroup members has an influence on face memory (Levin, 2000; Bodenhausen et al., 2003). When it comes to ingroup members, individuals may look for more specific facial differences than if they were to focus on broader characteristics in outgroup members (Levin, 2000; Hugenberg et al., 2007). Other work has shown that faces that are seen as more socially relevant may be better remembered (e.g., Wilson et al., 2014), which could help to explain why majority group perceivers

tend to show a larger CRE than minority group members. In this study, trustworthiness will be the driving factor for determining facial memory of ingroup and outgroup individuals.

Perceptions of Facial Trustworthiness

Race is of course not the only social dimension that impacts how people judge faces. There is a large literature, for instance, on the inferences that people make about others and their personality, based just on a quick glimpse. When it comes to making judgment based on facial appearance, conclusions can be made as quickly as 100ms (Willis and Todorov, 2006). Social groups are not the only factor influencing face memory, though. One facial characteristic that has recently been found to influence memory is the appearance of facial *trustworthiness* (Todorov et al., 2009; Rule et al., 2012). This body of work shows that people judge the trustworthiness of others quickly and with great consensus, despite the fact that trustworthiness judgments do not tend to be particularly accurate (Rule et al., 2013). As a result, encoding the faces of others occurs on the basis of certain basic facial traits; one of those basic traits happens to be trustworthiness, which essentially indexes another's presumed intentions (Todorov et al., 2010). Being perceived as untrustworthy can have severe consequences. Untrustworthy-looking targets are more likely to be presumed guilty for alleged crimes (Berry & Zebrowitz-McArthur, 1988), more likely to receive harsher sentences for hypothetical crimes (Porter et al., 2010), and even more likely to receive the death penalty relative to more trustworthy-looking targets who committed the same category of crime (Wilson & Rule, 2015).

A Memory Advantage for Untrustworthy Faces

Beyond these impacts on outcomes relevant to legal outcomes, having a trustworthy face can impact more everyday perceptions as well. Relevant to the present work, some research has actually shown that people more accurately remember untrustworthy faces than trustworthy faces (Rule et al., 2012). Other studies have shown that untrustworthy and unlikable faces yield a greater memory than trustworthy, likable faces (Oda, 1997; Mueller et al., 1988). These authors have posited that this occurs because we remember faces that are distinctive, whose appearances are atypical (Light et al., 1979; Vokey & Read, 1992). More directly relevant to facial trustworthiness, it is argued that untrustworthy faces may signal danger (Buchner et al., 2009; Cosmides, 1989; Suzuki & Suga, 2010).

The findings reported by Rule et al. (2012) are interesting because they do seem to confirm that there is a memory advantage for untrustworthy faces. However, they are also limited, in that they only tested memory for White faces with a mostly White participant sample. There remain open questions about whether such findings might be replicated in a task using faces of more than one race, and one in which non-White participants are recruited. It is plausible that the same memory advantage may be observed across both target and participant race. However, it also stands to reason that target race could moderate the impact of facial trustworthiness on face memory. Perhaps, for example, White participants might show even more highly elevated memory for untrustworthy faces (and/or worse memory for trustworthy faces) if those faces are Black. From a functional perspective, this could occur because it may be especially important to remember untrustworthy-looking outgroup targets.

However, there is some reason to believe that this may not occur. This thesis builds upon a previous unpublished study in which White participants showed a memory advantage for untrustworthy White faces, but no such advantage for untrustworthy Black faces (Wilson, in prep). In fact, these participants showed a non-significant tendency to more accurately remember trustworthy Black faces. In a separate pilot study, participants were asked to rate each face used in that memory study according to how distinctive it appeared to them. Consistent with the memory results, trustworthiness was negatively correlated with distinctiveness for White faces, but no such correlation occurred for Black faces. This raises the possibility that untrustworthy own-race faces appear distinctive, or "stick out," to perceivers, but that this does not occur for other-race faces. It is not yet clear why this occurs, but perhaps untrustworthy outgroup faces do not stick out if we expect outgroup members to be less trustworthy in general.

The Present Work: Race and Trustworthiness in Memory

The present work will build upon the unpublished data by using a larger sample of targets and participants, and crucially, it will also involve samples of both White and Black participants. This latter component is important, because it will allow for testing the generalizability of findings if the preliminary study is replicated. If only White participants are included, it will be unclear whether a successful replication means that perceivers *in general* do not show an untrustworthiness advantage. The inclusion of Black participants will allow for the full crossing of target and participant race, such that all participants view both ingroup and outgroup faces, and all targets are ingroup targets for some participants and outgroup targets for others.

As such, facial trustworthiness will be crossed with both target race and participant race in a 2 (Participant Race: White vs. Black) × 2 (Target Race: Ingroup vs. Outgroup) × 2 (Target Trustworthiness: Trustworthy vs. Untrustworthy) mixed design. Overall, it was predicted that there would be a main effect of facial trustworthiness, such that untrustworthy targets are recognized more accurately than trustworthy targets. However, critically, this main effect was expected to be qualified by interactions with the other IVs, such that the untrustworthiness advantage would weaken, eliminate, or even reverse for racial outgroup targets. Exploratory analyses were planned to test for the possibility that White and Black participants may show different patterns of memory for ingroup/outgroup trustworthy/untrustworthy faces. This work will provide an important update to the findings of Rule et al (2012) and will shed new light on the potential interactive effects of multiple social dimensions in face memory.

Method

Participants

Participants were 169 Montclair State University undergraduate students (132 female, 37 male; M_{age} = 20.0, *SD* = 2.6). One hundred twelve were non-Hispanic White and 57 were Black. An additional 59 participants completed the study but were not included in analyses because they reported a race other than White or Black. Recognition data of an additional 3 participants were not interpretable because they provided the same response for every target. This left 166 participants in primary analyses. The participants were recruited through the SONA program and consisted of undergraduate students enrolled in Psychology 101. Each participant provided informed consent before completing the study and received a debriefing form at the end.

Stimuli and Procedure

The study was conducted using Inquisit 4. Each participant sat in an individual room and completed the study with the door closed. The program allows stimuli to be displayed and the participant's responses to be recorded. After reading the instructions, participants underwent the main procedure.

Figure 1

Example Stimuli

Note. Trustworthy targets on left, untrustworthy targets on right.

There were 96 total male faces used in this study: 48 White and 48 Black (half trustworthy and half untrustworthy). Each stimulus was selected from a larger database of 200 (100 White, 100 Black) faces and rated by a separate sample of participants for trustworthiness. The 24 most and least trustworthy targets of each race were selected for inclusion in this study. In the learning phase, each participant viewed half of the total faces (48), one at a time, presented centrally onscreen for 3s each. The photo stimuli were displayed in a random order and participants were asked to simply pay attention to each face. There were not any instructions prompting them to remember any of the faces.

Once the learning phase concluded, the participants were directed to complete the Belief in a Dangerous World questionnaire (Altemeyer, 1988) dealing with race and social interactions. This served to show how the participants' belief of danger in the world influenced their recognition of trustworthy and untrustworthy stimuli and was administered as a Likert scale. Each response was to be answered on a scale of 1 to 7. This measure was a filler and responses were not analyzed.

Once the questionnaire was completed, the "recognition phase" of the study began. The participants were asked to complete a memory recognition task of the facial stimuli. Ninety-six stimuli were presented. Of these, 48 stimuli were "new" faces that were not seen before while the remaining 48 stimuli were "old" stimuli that were previously observed in the encoding phase. The participant was asked to recall if they had seen the face before or not. The directions read as follows:

Recognition: "Now that you have completed the questionnaires, it is time to move on to the RECOGNITION phase of the experiment. In the LEARNING phase of this study, you learned a series of faces. In this RECOGNITION phase, you will be shown a larger series of faces at the center of the screen, one face after another. Your job will be to decide whether each of the faces you see was among the faces that you learned during the LEARNING phase (i.e., an 'old' face that you've seen before) or whether this face is a 'new' face that was NOT in the LEARNING phase."

In the recognition phase, each stimulus stayed on screen until an old/new response was rendered via keystroke. The directions advised the participants not to spend too much time on any given stimuli. Once all 96 faces had been viewed and categorized as old or new, the participant was asked to answer basic demographic questions regarding age,

gender, and race. At this point, the study was finished and the research assistant was notified when the participant opened the door of the study room. All participants were debriefed.

Results

Preliminary Data Cleaning

First, all responses were coded to allow for signal detection analyses. This involved coding each response as either a Hit (correctly identifying a previously seen face) or a False Alarm (incorrectly indicating that a new face was previously seen). The other two types of responses, Misses and Correct Rejections, were not used for the analysis in this study. Hit and False Alarm proportions were converted to d', which is an overall measure of recognition sensitivity. For all analyses, Target Race was coded as Ingroup or Outgroup, such that White targets were coded as Ingroup for White participants and Outgroup for Black participants, and vice versa for Black targets.

Primary Analysis: Main Effects

For the primary hypothesis test, d' scores were subjected to a 2 (Target Trustworthiness: Trustworthy vs. Untrustworthy) $\times$ 2 (Target Race: Ingroup vs. Outgroup) $\times$ 2 (Participant Race: White vs. Black) mixed ANOVA, with repeated measures on the first two factors. First, as hypothesized, there was a significant main effect of target race, $F(1, 164) = 28.97$, $p < .01$, $\eta^2 = .025$, such that ingroup targets were recognized more accurately than outgroup targets. Next, consistent with previous research, there was a significant main effect of trustworthiness, $F(1, 164) = 70.99$, $p < .001$, $\eta^2 = .052$, such that untrustworthy targets were recognized more accurately than trustworthy targets. There was an unexpected main effect of the participant race, $F(1,164)$

= 5.92, $p = .02$, $\eta^2 = .018$, such that Black participants showed better face memory than White participants. These main effects were qualified by higher-order interactions, which are described in the next section.

Primary Analysis: Interactions

There were three significant interactions, which can be seen in Figures 2 and 3. First, there was a two-way interaction between Target Race and Participant Race, $F(1, 164) = 32.49$, $p < .001$, $\eta^2 = .028$, such that White participants, but not Black participants, showed a significant CRE. Critically, there was also a significant interaction between Trustworthiness and Target race, $F(1, 164) = 7.35$, $p < .001$, $\eta^2 = .005$. As can be seen in Figure 2, this interaction is indicative of support for the primary hypothesis, such that the untrustworthiness advantage was significantly stronger for ingroup targets than outgroup targets.

Figure 2

Recognition Sensitivity across All Participants

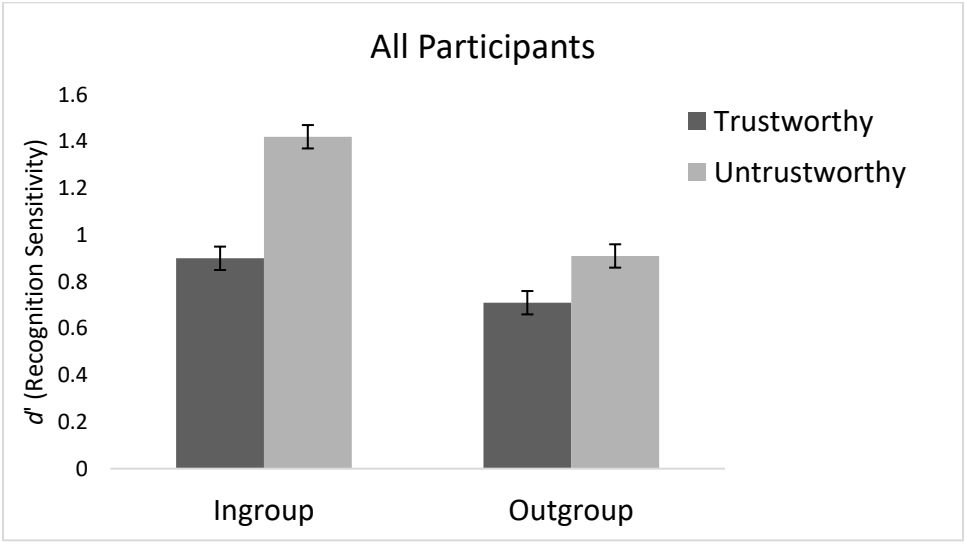

Note. Error bars represent +/- 1 standard error.

Breaking down this two-way interaction, untrustworthy ingroup targets (d' = 1.41, SD = 0.78) were recognized much more accurately than trustworthy ingroup targets (d' = 0.90, SD = 0.72), $t(166) = 8.72$, $p < .001$, Cohen's d = 0.68. Untrustworthy outgroup targets (d' = 0.91, SD = 0.76) were also recognized more accurately than trustworthy outgroup targets ((d' = 0.71, SD = 0.67), $t(166) = 3.56$, $p < .001$, but the effect was smaller, Cohen's d = 0.28.

Finally, these two-way interactions were qualified by a significant three-way interaction, $F(1, 164) = 6.74$, $p = .01$, $\eta^2 = .005$. This will be broken down further by investigating the effects separately by participant race. As can be seen in Figure 3, the moderating effect of target race on the untrustworthiness advantage only occurred among White participants. For Black participants, untrustworthy targets (d' = 1.33, SD = 0.56) were recognized more accurately than trustworthy targets (d' = 0.93, SD = 0.53), F(1, 53) = 25.38, $p < .001$, $\eta^2 = .32$, and there was no interaction between trustworthiness and target race, F(1, 53) = .01, $p = .94$, $\eta^2 = <.001$. Among White participants, untrustworthy targets (d' = 1.08, SD = 0.66) were again recognized more accurately than trustworthy targets (d' = 0.74, SD = 0.59), F(1, 111) = 51.08, $p < .001$, $\eta^2 = .32$. However, this main effect was qualified by a two-way interaction between trustworthiness and target race, F(1, 111) = 21.30, $p < .001$, $\eta^2 = .16$, such that the untrustworthiness advantage occurred for racial ingroup targets, $t(111) = 7.99$, $p < .001$, Cohen's d = 0.75, but not racial outgroup targets, $t(111) = 1.67$, p =.10, Cohen's d = 0.16. Overall, these results show that, for White participants, the untrustworthiness advantage in memory may be limited to ingroup faces. However, for Black participants, the untrustworthiness advantage in memory may occur regardless of target group.

Figure 3

Recognition Sensitivity Plotted Separately by Participant Race

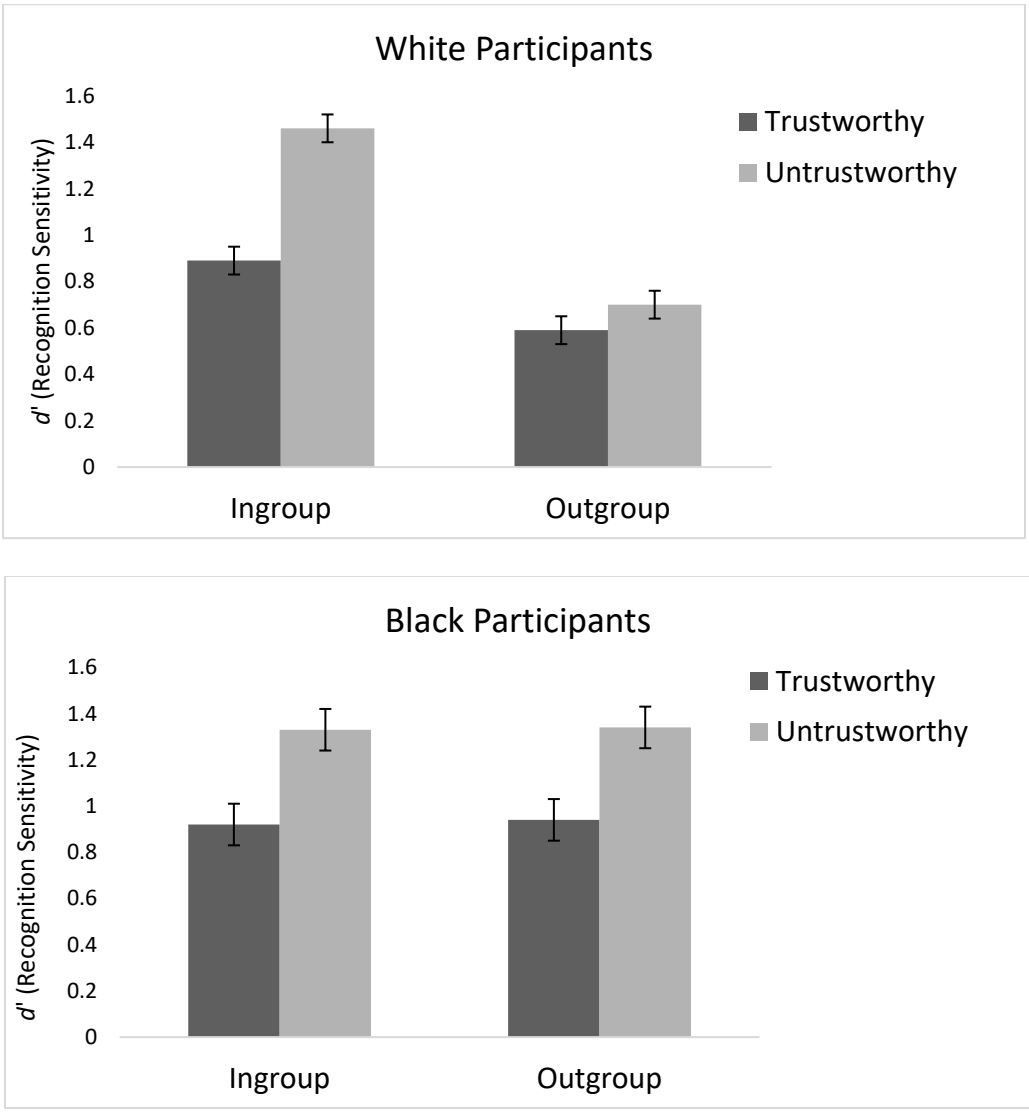

Note. Error bars represent +/- 1 standard error.

Discussion

A number of interesting findings emerged in this study. First, confirming previous findings, participants overall did display the CRE in face memory. However, whether or not the CRE occurred depended upon participant race, as only White participants showed this effect. The fact that Black participants did not show the CRE is not entirely surprising, and is in fact consistent with some previous empirical (Feinman & Entwisle, 1976) and meta-analytical (Meissner & Brigham, 2001) research.

Furthermore, participants overall showed an untrustworthiness advantage, such that they remembered untrustworthy targets more accurately than trustworthy targets. This untrustworthiness advantage was moderated in a couple of key ways. First, it was moderated by target group. Consistent with the primary hypothesis, and confirming pilot research, participants showed a larger untrustworthiness advantage for ingroup targets than for outgroup targets. However, a significant three-way interaction showed that this finding was even more nuanced. As in previous research, White participants showed a significant untrustworthiness advantage for racial ingroup targets, but not racial outgroup targets, who were recognized rather poorly regardless of trustworthiness. However, Black participants showed a strong untrustworthiness advantage regardless of target racial group.

These findings have a number of potentially interesting implications. First, they add to a large literature confirming that people do tend to show better memory for ingroup faces than outgroup faces, but that this may occur more strongly for racial majority group members. This finding is consistent with both perceptual expertise (Rhodes et al. 1989) and social cognitive (e.g., Hugenberg et al., 2007) accounts of the

CRE. Minority perceivers are likely to have more experience viewing outgroup faces than majority perceivers, and they also may be more likely to judge outgroup faces as socially relevant or important, especially if those outgroup targets belong to a majority group.

There are also implications regarding the untrustworthiness advantage. As found in Rule et al. (2012), untrustworthy faces yielded better recall than trustworthy faces. Taken alone, this finding supports the ideas put forth by those authors - that untrustworthy faces may be seen as particularly high in ecological value (Buchner et al., 2009; Suzuki & Suga, 2010) because they signal a danger that should be attended to and perhaps acted upon. However, emerging data from our lab had suggested that the previously observed memory advantage for untrustworthy faces (Rule et al., 2012) may not be robust across different target racial groups. The present work aimed to be a systematic test of the generalizability of this finding for two racial groups (White and Black), and importantly, for both White and Black participants. Indeed, in the present work, not all participant groups showed this tendency for all target groups. Specifically, White participants did not show particularly accurate memory for outgroup (Black) untrustworthy targets. From an ecological perspective that focuses on functional value, this seems potentially surprising. Outgroup members, especially those who belong to a group about which many threat-related stereotypes exist (Cottrell & Neuberg, 2005), may be particularly relevant as objects of attention if they have untrustworthy appearances. On the other hand, perhaps it is the case that the trustworthiness of one's facial appearance does not impact distinctiveness and memorability if the person displaying that appearance

belongs to a group that is seen as generally threatening. In the present study, this latter possibility seems to have been supported.

More research will need to be done in order to uncover the underlying psychological processes involved in producing these memory data. This work may be needed to disentangle a functional account from distinctiveness-based accounts, and should perhaps consider the possible role of expectancy violation in determining who is memorable and who is not. The cause for the trend of White Participants' responses could, for example, be due to the false sense of familiarity that comes with the Black outgroup faces. If one group tends to more accurately recall faces of their own racial ingroup and not so for racial outgroups, it could come with potentially serious consequences. When applied to real-life circumstances, this false familiarity of generalizing racial outgroups occurring when needing to accurately recall a criminal in a police line-up could result in a wrongful conviction of an innocent person.

Strengths, Weaknesses, and Future Directions

This research has a number of strengths. First, it provides a much-needed replication of previous work. Next, it extended upon that work in two crucial directions, by virtue of including both participants and targets of two different racial groups. This approach allows the present study to merge two previously distinct phenomena (the CRE and the untrustworthiness advantage in memory) into one, while accounting for potentially interactive effects of the central independent variables. The inclusion of Black participants in addition to White is particularly valuable, especially given that the memory phenomenon in question is often referred to generally as the "Cross Race Effect" or "Own Race Bias." For such a phenomenon to truly be cross-race, researchers cannot

simply recruit participants of one or another racial group. Similarly, if researchers are to assume that people *in general* show better memory for untrustworthy faces, it is important to test whether this occurs narrowly for same-race faces or indeed occurs more broadly. The answer to this question is complicated, and we would not know this without having taken the present approach.

Unfortunately, this work does have a number of weaknesses. For one, it was not possible to recruit an equivalent number of Black and White participants within the timeframe of the study. Next, due to the nature of the participant pool, women made up a disproportionate share of participants. Further, the mean age in this study was 20. This study was concerned with testing the generalizability of phenomena observed in previous research, but this test does have limits because of the sample characteristics. In the future, this work should be replicated in a participant pool that is more representative of the general population. Finally, this study only tested participants on recall of male faces. It would be important to replicate these findings with female targets in the future.

Another interesting future avenue would be to include targets and participants of other races or ethnicities. Doing so would allow researchers to better identify the conditions that may be necessary for an untrustworthiness advantage to emerge. Perhaps this phenomenon does not occur for any outgroup about which threat stereotypes are common, or perhaps the phenomenon hinges more upon whether participants and targets are members of numerical minority and majority groups. These questions can be answered in future work.

Conclusion

This study can shed light on the potential underlying factors behind racial and facial appearance-based biases in face memory, and on how face memory may be affected by intersections of stereotypes, appearance, and social identity. Who we remember may be driven by all manner of factors - social relevance and power, fear, or the sense that a person is a warm and welcoming affiliation partner. In such a diverse world, it is critical to understand all of the factors that impact how we interact with one another. The present research helps to move us further along the road toward understanding some of these factors, and it advances our knowledge in multiple meaningful ways. As is clear from this work, social perception is endlessly complex, and any attempt to uncover general principles underlying psychological phenomena must be undertaken with nuance and care.

References

Altemeyer, B. (1988). Enemies of freedom: Understanding right wing authoritarianism. *San Francisco: Jossey-Bass*.

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.

Berry, D. S., & Zebrowitz-McArthur, L. (1988). What's in a face? Facial maturity and the attribution of legal responsibility. *Personality and Social Psychology Bulletin*, *14*(1), 23-33.

Bodenhausen, G. V., Macrae, C. N., & Hugenberg, K. (2003). Social cognition.

Bothwell, R. K., Brigham, J. C., & Malpass, R. S. (1989). Cross-racial identification. *Personality and Social Psychology Bulletin*, *15*(1), 19-25.

Brigham, J. C., & Malpass, R. S. (1985). The role of experience and contact in the recognition of faces of own-and other-race persons. *Journal of social issues*, *41*(3), 139-155.

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*(3), 212-224.

Chiroro, P., & Valentine, T. (1995). An investigation of the contact hypothesis of the own-race bias in face recognition. *The Quarterly Journal of Experimental Psychology*, *48*(4), 879-894.

Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition*, *31*(3), 187-

276.

Cottrell, C. A., & Neuberg, S. L. (2005). Different emotional reactions to different groups: a sociofunctional threat-based approach to" prejudice". *Journal of personality and social psychology*, *88*(5), 770.

Feinman, S., & Entwisle, D. R. (1976). Children's ability to recognize other children's faces. *Child development*, 506-510.

Hastie, R., & Kumar, P. A. (1979). Person memory: Personality traits as organizing principles in memory for behaviors. *Journal of Personality and Social Psychology*, *37*(1), 25.

Hugenberg, K., Miller, J., & Claypool, H. M. (2007). Categorization and individuation in the cross-race recognition deficit: Toward a solution to an insidious problem. *Journal of Experimental Social Psychology*, *15*(2), 334-340.

Hugenberg, K., Wilson, J. P., See, P. E., & Young, S. G. (2013). Towards a synthetic model of own group biases in face memory. *Visual Cognition*, *21*(9-10), 1392-1417.

Levin, D. T. (2000). Race as a visual feature: using visual search and perceptual discrimination tasks to understand face categories and the cross-race recognition deficit. *Journal of Experimental Psychology: General*, *129*(4), 559.

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*(3), 212.

Mason, M., Hood, B., & Macrae, C. N. (2004). Look into my eyes: Gaze direction and person memory. *Memory*, *12*(5), 637-643.

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, and Law*, *7*(1), 3.

Mueller, J. H., Thompson, W. B., & Vogel, J. M. (1988). Perceived honesty and face memory. *Personality and Social Psychology Bulletin*, *14*(1), 114-124.

Oda, R. (1997). Biased face recognition in the prisoner's dilemma game. *Evolution and Human Behavior*, *18*(5), 309-315.

Porter, S., ten Brinke, L., & Gustaw, C. (2010). Dangerous decisions: The impact of first impressions of trustworthiness on the evaluation of legal evidence and defendant culpability. *Psychology, Crime & Law*, *16*(6), 477-491.

Rhodes, G., Brake, S., Taylor, K., & Tan, S. (1989). Expertise and configural coding in face recognition. *British journal of psychology*, *80*(3), 313-331.

Rule, N. O., Krendl, A. C., Ivcevic, Z., & Ambady, N. (2013). Accuracy and consensus in judgments of trustworthiness from faces: behavioral and neural correlates. *Journal of personality and social psychology*, *104*(3), 409.

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

Suzuki, A., & Suga, S. (2010). Enhanced memory for the wolf in sheep's clothing:: Facial trustworthiness modulates face-trait associative memory. *Cognition*, *117*(2), 224-229.

Tanaka, J. W., Kiefer, M., & Bukach, C. M. (2004). A holistic account of the own-race effect in face recognition: Evidence from a cross-cultural study. *Cognition*, *93*(1), B1-B9.

Todorov, A., Loehr, V., & Oosterhof, N. N. (2010). The obligatory nature of holistic processing of faces in social judgments. *Perception, 39*(4), 514-532.

Todorov, A., Pakrashi, M., & Oosterhof, N. N. (2009). Evaluating faces on trustworthiness after minimal time exposure. *Social Cognition, 27*(6), 813-833.

Vokey, J. R., & Read, J. D. (1992). Familiarity, memorability, and the effect of typicality on the recognition of faces. Memory & Cognition, 20(3), 291-302.

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

Wilson, J. P., & Rule, N. O. (2015). Facial trustworthiness predicts extreme criminal-sentencing outcomes. *Psychological science*, *26*(8), 1325-1331.

Wilson, J. P., See, P. E., Bernstein, M. J., Hugenberg, K., & Chartier, C. (2014). Differences in anticipated interaction drive own group biases in face memory. *PloS one*, *9*(3), e90668.