PLAGIARISM DECLARATION

PLAGIARISM

This means that you present substantial portions or elements of another's work, ideas or data as your own, even if the original author is cited occasionally. A signed photocopy or other copy of the Declaration below must accompany every piece of work that you hand in.

DECLARATION

1. I know that Plagiarism is wrong. Plagiarism is to use another's work and pretend that it is one's own.

2. I have used the American Psychological Association formatting for citation and referencing. Each significant contribution to, and quotation in, this essay/report/project from the work or works, of other people has been attributed, cited and referenced.

3. This essay/report/project is my own work.

4. I have not allowed, and will not allow anyone to copy my work with the intention of passing it off as his or her own work.

NAME: Daniel Derbyshire

SIGNATURE: DDerbyshire

STUDENT NUMBER: DRBDAN001

Investigating the effects of race-based exclusion and prejudice on the own-race bias in facial recognition

Daniel Derbyshire (DRBDAN001)

ACSENT Labarotary
Department of Psychology
University of Cape Town

Supervisor: Professor Colin Tredoux

Word count:
Main body**:** 7996 words
Abstract: 216 words

Abstract

The own-race bias (ORB) in facial recognition refers to the finding that perceivers in face recognition tasks demonstrate better recognition accuracy for faces of their racial in-group compared to faces of a different race. The effect has been reliably demonstrated, but the mechanisms underlying it are still debated. I investigated the role of prejudice in producing the ORB when participants have been excluded or included based on race using the *Cyberball* paradigm to manipulate social exclusion. The face recognition task was built into the *Cyberball* as it consisted of the encoding phase of the face recognition task. I used a mixed factorial design to investigate the effects of social exclusion/inclusion and race of participant on cross-race recognition accuracy. Implicit prejudice, measured by the Implicit Association Task, was included as a covariate. I analysed data on a final sample of 92 participants, 59 white and 33 black. Black participants demonstrated a significant ORB but white participants did not. Social exclusion and inclusion was not significantly associated with the ORB. The effect of implicit prejudice, on the ORB was also not significant. There was a significant effect of condition, where participants in the control group demonstrated significantly better recognition accuracy than the experimental groups. Overall, these results suggest that implicit prejudice is not a predictor of the ORB.

Key words: Cyberball; exclusion; face recognition; own-race bias; prejudice; Race

The own-race bias (ORB), or cross-race effect (CRE), is the tendency for perceivers, in face recognition tasks, to recognise same race (SR) faces more accurately than cross-race (CR) faces (Young & Hugenberg, 2012). The ORB has been reliably found and replicated across a number of psychological studies investigating the accuracy of facial recognition. However, the underlying mechanisms behind the ORB are still largely debated. One possible factor that has received little attention and mixed support in the ORB literature is implicit prejudice toward different racial groups, which refers to unconscious biases or negative attitudes toward those groups (Qian, et al., 2017; Young, Hugenberg, Bernstein, & Sacco, 2012). Implicit prejudice may play an important role in how CR faces are categorised or represented in the mind of perceivers, which could influence subsequent processing of these faces. Additionally, negative contact, especially social exclusion, which is associated with group identification and implicit prejudice, may play an important role in explaining the ORB, as increased negative contact is associated with increased category salience (Paolini, Harwood, & Rubin, 2010). In turn, increased category or group salience may affect the way SR and CR faces are processed.

Explanations of the ORB are broadly divided into perceptual expertise models, and social cognitive models. Perceptual expertise models propose that a lack of contact, or expertise, with CR faces results in differences in SR and CR face recognition (Young et al., 2012). A central concept in perceptual expertise perspectives is the role of interracial contact in improving CR recognition; increased contact with other races is thought to facilitate better CR recognition through 'practice' or familiarity with those faces, which translates into better processing of CR faces and an increased ability to distinguish between different CR faces (Shriver, Young, Hugenberg, Bernstein, & Lanter, 2008). Alternate explanations for the ORB draw on social cognitive theories, specifically the idea that out-group members are processed differently to those belonging to the perceivers' in-group. According to social cognitive accounts, in-group members are processed in an individuating manner through focusing on facial features that can be used to differentiate one in-group member from another. Conversely, out-group members are processed more categorically by focusing on features that indicate broad group membership of the face, such as race or gender (Young et al., 2012).

The ORB has been widely demonstrated across studies in many different countries, but is most frequently studied using white and black faces and participants. White perceivers, in

particular, have been shown to demonstrate a significant bias toward white faces presented during face processing tasks when compared to recognition of black faces. This effect has been found many contexts where white people either make up the majority of the population (UK and USA) or have more social power (South Africa). Some studies have even shown that in contexts where white people have more social power, often through financial and political dominance, black participants show similar recognition for black and white faces, regardless of whether white people form part of a majority or minority group (Shriver & Hugenberg, 2010; Wright, Boyd, & Tredoux, 2003). Similar results have been found for white participants' recognition of South Asian faces in the UK (Walker & Hewstone, 2008). These results indicate that the ORB can be influenced by social and contextual factors.

Support for social cognitive and perceptual expertise accounts is mixed. For example, while a number of studies conducted in the US, UK, and South Africa have found high self-reported interracial contact to be associated with a reduced ORB, interracial contact accounts for very little variance in the ORB (Meissner & Brigham, 2001). In terms of social cognitive models, investigations of the mechanisms underlying this differential processing of SR and CR faces have produced varying results (Young et al., 2012). Therefore, a third, integrative model has been put forward to attempt to explain these differences in findings. The Categorisation-Individuation model (CIM) explains CR deficits as a combination of perceiver expertise and motivation to individuate CR faces (Young & Hugenberg, 2012; Young et al., 2012). According to the model, increased motivation to process characteristics that distinguish one CR face from another interacts with increasing experience or familiarity to reduce CR recognition deficits (Young & Hugenberg, 2012). Experiments designed to test this model have shown some support for this integrative explanation of the ORB but more work is needed, as there have been some failures to replicate these results (Lebrecht, Pierce, Tarr, & Tanaka, 2009; McGugin, Tanaka, Lebrecht, Tarr, & Gauthier, 2011; Young & Hugenberg, 2012). Considering the lack of consistent findings to explain the ORB, other factors such as prejudice remain potentially important predictors of CR recognition deficits.

Prejudice and the ORB

The specific role of prejudice has been explained and investigated using all three models outlined above, but has received mixed results in the few studies designed to directly examine its influence on and relationship with the ORB. While previous investigations have examined the

role of prejudice through self report measures and found little evidence in support of its role (Brigham & Barkowitz, 1978; Ferguson, Rhodes, Lee, & Sriram, 2001), measures of implicit prejudice may be better suited for studies of the ORB as they assess peoples' automatic, unconscious biases (Ferguson et al., 2001; Hugenberg & Bodenhausen, 2003).

Recent evidence suggests that implicit measures of prejudice are associated with CR deficits, where participants that measured higher on measures of implicit prejudice were shown to have poorer CR recognition (Ma, Yang, & Han, 2011; Walker & Hewstone, 2008; Young et al., 2012). However, this finding is inconsistent, as other studies have not found a significant relationship between implicit prejudice and the CRE (Ferguson et al., 2001; Young et al., 2012). Further evidence for the role of prejudice comes from experiments that have utilised perceptual individuation training. Perceptual individuation training teaches participants to individuate CR faces, or process CR faces at an individual level rather than a categorical level. In the few studies that have used this method, individuation training was found to significantly reduce the ORB (Lebrecht et al., 2009; Qian et al., 2017). These results illustrate the possible importance of racial prejudice or implicit bias in the CRE, as teaching individuals to regard CR targets as individuals rather than merely group members can lead to a decrease in the both the ORB and implicit prejudice.

However, consistent with perceptual expertise theories, interracial contact may still also still play an important role in the ORB, but evidence suggests that negative encounters with CR individuals may contribute more to the ORB than positive encounters or the quantity of CR encounters (Barlow et al., 2012). Negative encounters may result in prejudiced or stereotyped representations of CR individuals. These representations can result in CR recognition deficits through homogenising effects that work to categorise them CR individuals as out-group members (Barlow et al., 2012).

Negative contact, social exclusion and the ORB

Following findings that negative contact may play a role important in producing the CRE and implicit prejudice, social exclusion represents another factor that could contribute to cross-race recognition differences, and may interact with prejudice to help explain the CRE. From a developmental perspective, social exclusion represents a powerful form of negative contact that has been demonstrated to be important in shaping early childhood prejudice (Rutland & Killen, 2015). In-group identification is an important factor in leading to social exclusion and prejudice,

as children are more likely to show prejudice toward out-group members and exclude them when they show greater identification with their in-group (Rutland & Killen, 2015). Consistent with these developmental perspectives, social exclusion has been found to lead some participants to identify more with their in-group, possibly in an attempt to protect against the negative effects of exclusion (Knowles & Gardner, 2008). Consequently, their increased identification with an in-group can result in changes in the way participants perceive, categorise, and attend to out-group members, and exclusion has been found to lead to increased categorisation of out-group members, which may lead participants to pay less attention to individuating information for out-group members (Gaither, Pauker, Slepian, & Sommers, 2016).

In this way, social exclusion may produce an ORB when participants are required to remember SR and CR faces in the context of social exclusion and increased in-group identification. Prejudice may be an important factor in this relationship, as those high in prejudice may display higher in-group favouritism and more biased or preferential encoding of racial in-group members (Guinote, Willis, & Martellotta, 2010). The experience of negative contact created through exclusion may also therefore activate prejudiced representations of out-group members that facilitate categorical processing rather than individuation (Barlow et al., 2012). Further evidence for the hypothesis that negative contact is important producing the CRE comes from findings by Paolini et al., (2010) that suggest that negative intergroup contact leads to increased category salience, meaning intergroup differences are emphasised and membership to specific groups is more salient (the valence-salience effect). Therefore, negative inter-group contact between different racial groups may result in increased group salience, which in turn results in differential encoding of in-group and out-group members.

Aims, and hypotheses

Following evidence for the role of social exclusion and negative contact in producing the ORB, and the relationship between social exclusion and prejudice, I used the game *Cyberball*, a computer-based ball tossing game to simulate social exclusion or inclusion (Williams, Cheung, & Choi, 2000). The game has been successfully demonstrated as a method for increasing participants' group identification, especially in groups that are invariant or perceived as biologically based, such as race (Knowles & Gardner, 2008). By using *Cyberball* to manipulate social exclusion and induce feelings of identification with participants self-identified racial group, I hoped to further investigate the role of prejudice in producing the CRE in a way that will

differentiate the performance of high and low prejudice participants (Gaither et al., 2016; Sacco, Wirth, Hugenberg, Chen, & Williams, 2011).

1. Race of participant will affect their recognition of CR vs. SR faces, where participants will be better at recognising SR faces when compared to CR faces.
2. Social exclusion and inclusion will result in an increase in the CRE relative to the control condition as group differences are made salient.
3. Implicit prejudice will be associated with the CRE, where participants' with higher prejudice scores will show poorer CR recognition performance.

Methods

Design and Setting

I used a 3X2X2 mixed design to investigate the effects of race-based social exclusion on the CRE. The independent variables were condition (exclusion vs. inclusion vs. control), race of participant (black vs. white). The within-subjects factor was the CRE (SR face recognition vs. CR face recognition). Implicit prejudice was measured as a continuous variable so I included it as a covariate in my analysis. The study included three groups or conditions, two experimental and one control. The entire study was run through the online survey platform Qualtrics. The study took place in the ACSENT laboratory in the UCT psychology department and was also distributed as an online survey.

Participants

I conducted an a priori power analysis using G*Power to calculate the sample size required for the study. To detect an effect of .25 with a power of .80 the sample size was calculated at 108. I therefore attempted to recruit 120 participants so there were 40 assigned to each condition. I used conventions set by Cohen, as I could not find any studies sufficiently similar to this one (Cohen, 1988). I collected a total of 132 responses from participants who identified as black, white, coloured, Indian/Asian, and other. No other demographic information was collected about participants. However, while anyone was allowed to participate, only data from black and white participants were included in the final analyses, as only the study was only investigating the ORB in white and black participants. A total of 92 black and white participants were included in the final analyses. There were 59 white and 33 black participants. Therefore, I

did not reach the desired sample size for black and white participants, reducing the overall statistical power.

Participants were recruited through convenience sampling methods, and data was collected using a mix of in-person collection and through the distribution of the study as an online survey. UCT undergraduate Psychology students were recruited as participants in the study through the Psychology department's Student Research Participation Program (SRPP) and through announcements appealing for responses posted on undergraduate Vula sites. Additionally, I shared the experiment/survey on social media and asked friends and classmates to share it with anyone who would be willing to participate. Participants with a current or previous diagnosis of a disorder included in the DSM were excluded from participation due to the negative emotional effects associated with social exclusion (Van Harmelen et al., 2014). There were no other exclusion criteria. The study took 30 to 45 minutes to complete depending on the condition participants were assigned to. Participants recruited through SRPP were awarded 2 points for their participation.

Instruments

Sociodemographic question. Participants were only asked to report their self-described race according to the racial categories used in South Africa.

Implicit prejudice. To assess implicit prejudice, I used the Implicit Association Task (IAT) (Greenwald, McGhee, & Schwartz, 1998). The IAT is a computer-based procedure that measures the association between two categories (e.g. white and black people) and different attributes (e.g. pleasant and unpleasant). Participants press a single key on a computer keyboard to respond to a category or an attribute presented one at a time (Teige-Mocigemba, Klauer, & Sherman, 2010). Typically, the IAT consists of 7 blocks with 20-40 trials per block (Greenwald, Nosek, & Banaji, 2003). It is assumed that when a category and attribute that are perceived to be highly associated by the participant share a response key, their response latencies will be shorter than when a category-attribute response key are seen as incompatible (Scroggins et al., 2016). The IAT has been found to be a reliable measure of implicit prejudice, with α reported between .7 and .9 (Teige-Mocigemba et al., 2010).

The IAT was created using the website Iatgen (https://applibs.shinyapps.io/iatui2/; Carpenter et al., 2018). Target A stimuli were six greyscale white faces and target B stimuli were six greyscale black faces. The faces were artifical faces created using Id, a program that creates

composite faces from a database of photos. Attribute A was six pleasant or positive words (cheer, enjoy, friend, happy, love, excitement), and attribute B was six unpleasant words (damage, gloom, awful, hurt, rotten, selfish). In compatible trials, white faces were paired with pleasant attributes, and in incompatible trials white faces were paired with negative stimuli, and vice versa for black faces. Therefore, faster responses by white participants during compatible trials and slower responses during incompatible trials resulted in increased implicit prejudice, as indicated by higher IAT scores. For black participants, faster responding to incompatible trials indicated increased implicit prejudice. The IAT was integrated into Qualtrics and consisted of 4 different permutations where the order of compatible and incompatible trials was varied. Participants were randomly assigned to one of the four permutations within Qualtrics.

Social exclusion. I used the *Cyberball* paradigm to manipulate social exclusion for the two experimental groups (Hartgerink, Van Beest, Wicherts, & Williams, 2015; Williams et al., 2000). The paradigm has been successfully used in hundreds of psychological studies to investigate the effects of social exclusion (Hartgerink, et al., 2015). *Cyberball* is a computer game where participants throw a virtual ball around with up to 8 other virtual players, who participants believe are real people. The researcher is able to manipulate social exclusion by programming how often the virtual players throw the ball to the participant (Williams, Yeager, Cheung, & Choi, 2012). In a typical inclusion condition, the other "players" consistently throw participant the ball during the game. In the exclusion condition, the participant is thrown the ball once or twice initially, but is then ignored for the rest of the game while the other "players" pass the ball to each other (Bernstein, Sacco, Young, & Hugenberg, 2014).

To manipulate race-based exclusion or inclusion, the in-game avatars consisted of faces of different races to create the idea that they were players of different races. In the exclusion condition, CR players excluded participants and other players of the same race of the participant. In the inclusion condition, players of the same race as the participant threw the ball only to the participant and other SR players. I created three different exclusion games and three different inclusion games, each with a unique throw schedule, unique set of faces, and different player positions (see Appendix A for in-game screenshots). Participants in the control condition did not play *Cyberball,* so there was no manipulation of social exclusion.

Face recognition task. The recognition task used in the study was a forced choice *yes/no* recognition task where 24 target faces were presented during encoding. Half of the faces were

black and half were white. During the recognition phase, the 24 target faces were presented amongst 24 distractor faces or foils (12 black and 12 white). In the experimental conditions, the face recognition task was integrated into the *Cyberball* games, with the consisting of the encoding phase of the recognition task. Target faces were displayed as the avatars of the players in the game. Each participant played three games of *Cyberball*, with eight different faces presented per game. Participants therefore encoded 24 target faces (12 black and 12 white). The order of games was randomised within Qualtrics to vary the order of presentation of the faces. During the recognition phase, the 48 faces were presented one at a time. Participants had to respond by selecting that "yes" they had seen a face previously or "no" they had not seen it previously. The order of presentation of faces during the recognition phases was also randomised. In the control condition, target faces were presented one at a time and displayed for 3 seconds (Bernstein et al., 2014; Meissner & Brigham, 2001; Young & Hugenberg, 2012). The order of presentation during encoding was also randomised. The recognition phase was the same as for the experimental groups.

The face stimuli consisted of 48 composite faces created using the program Id. The composites were created using a database of real photos using facial features from the photos. I did this so no real people were implicated in the game, as I did not want participants to recognise someone as a person they believed excluded them. I edited each face and added a white t-shirt to each picture in Photoshop to make them look as realistic as possible. The pictures were all the same size at encoding in the experimental conditions, as I could not change the size of the picture in *Cyberball*. In the recognition phase and in the encoding phase for the control condition all the faces were sized 149mm x 200mm. I varied the pictures from encoding to recognition by presenting the faces in colour at encoding but in greyscale at recognition, so participants did not simply remember the picture.

Manipulation check. I included two items as a manipulation check following the games of *Cyberball* (Bernstein et al., 2014). The items were "I was excluded" (item 1) and "I was ignored" (item 2), which participants respond to on a 7-point Likert scale from 1 (Strongly disagree) to 7 (Strongly agree). Higher responses during the exclusion condition and lower responses during the inclusion condition indicated that the manipulation was successful.

Affective prejudice. I used an affective prejudice scale (Appendix C) as a measure of explicit prejudice to examine participants' self-reported feelings toward members of the opposite

race (Finchilescu, 2010; Zanna, 1994). The scale consists of six items scored on a scale scored from 1 (Strongly disagree) to 7 (Strongly agree). Participants expressed their affective responses toward out-group members (black or white) by responding to 6 pairs of bipolar adjectives on the 7-point scale (Finchilescu, 2010). Higher scores indicate stronger negative affective or emotional feelings toward racial out-group members (Tredoux & Finchilescu, 2010). The scale has been used previously in a South Africa, and has shown high internal consistency, $\alpha = .87$ for black participants and .90 for white participants (Tredoux & Finchilescu, 2010).

Group identification. I included a group identification scale (Appendix B) adapted from Mackie, Devos, and Smith (2000). The scale provided an estimate of participants' identification with their racial in-group. The scale included four items that participants responded to for both their racial in-group and out-group. The statements were responded to on a 7-point Likert scale from 1 (Strongly disagree) to 7 (Strongly agree) (Mackie et al., 2000). Higher scores indicated higher feelings of identification or belonging with a racial group. The scale was demonstrated to be a reliable measure of group identification, $\alpha = .89$ for the in-group and .88 for the out-group (Mackie et al., 2000).

Intergroup emotion. I also included a measure of intergroup emotion (Appendix C) using items developed by Mackie et al., (2000). This scale measured participants' emotional responses to out-group members. The scale consisted of four items measuring anger and four items measuring fear toward racial out-groups. Participants responded will state their agreement to the extent that out-group members make them feel angry, displeased, irritated, furious, worried, anxious, afraid, or fearful on a scale from 1 (not at all) to 7 (extremely) (Mackie et al., 2000). Higher scores indicated higher negative emotions toward out-group members, namely increased anger toward and fear of out-group members. The items were found to have high internal consistency in two separate studies ($\alpha = .87$ and $\alpha = 88$ for the anger items; $\alpha = 89$ and $\alpha = .88$ for the fear items).

Procedure

All procedures were completed on a computer using Qualtrics. I created three separate surveys in Qualtrics (one for each condition). Participants in the lab were randomly assigned to a computer with one of the surveys open, depending on the condition they were assigned to. Participants who were recruited through social media or the undergraduate Vula sites to

complete the experiment as an online survey were provided a link that redirected them randomly to one of the three surveys.

For in-person data collection, I welcomed participants to the ACSENT lab, showed them to the computer they would be using, and explained that everything would be conducted on the computer. For all participants, informed consent was the first page of the survey. I instructed all participants to read through the informed consent form and begin when they were ready. On the second page, participants selected their race.

Group 1 (Exclusion condition). Following the informed consent page and race selection page, participants played three consecutive games of *Cyberball* where CR players excluded them and they encountered the target faces. Each game lasted approximately 2 minutes 30 seconds. They then completed the manipulation check. Participants then completed a 5-minute filler task where they had to find anagrams for a set of words. Following the filler task, participants completed the recognition phase of the face recognition task, where they viewed 48 faces (24 targets and 24 distractors). After the recognition phase, they completed the IAT, followed by the self-report measures. They first completed the group identification measure, followed by intergroup emotion, and finally the affective prejudice measure. Participants then played another three games of *Cyberball*, which were "include all" games. In these games, all players received the ball, including the participant. This set of games was included to attempt to counter the negative effects of exclusion. Participants then completed another IAT, followed by the debriefing form. The exclusion condition took approximately 45 minutes to complete.

Group 2 (Inclusion condition). Group 2's procedure was the same as Group 1, except that Group 2 only played the first set of *Cyberball* games during which SR players included them. The order of tasks was exactly the same as in group 1. In Group 2, however, participants saw the debriefing form directly after the self-report measures. The inclusion condition took approximately 30 minutes to complete.

Group 3 (Control condition). Participants in the control group did not undergo any exclusion or inclusion manipulation. After the informed consent form, they took part in the encoding phase during which faces they passively viewed target faces presented one at a time for 3 seconds. Following the encoding phase, they completed the same filler task. All procedures after the filler task were identical to group 2. The control condition took approximately 20 to 30 minutes to complete.

At the end, participants read through and signed the debriefing form on the computer. I provided all my contact information in the debriefing form if they had any questions or problems. For participants who took part in the study in the lab, I checked that they were fine to leave following the experiment and offered to answer any questions they may have.

Data analysis

Face recognition scores. Face recognition scores were analysed using Signal Detection Theory (SDT). SDT provides an estimate of participants' recognition accuracy (*d')* (Stanislaw & Todorov, 1999). *D'* provides an estimate of a participant's ability to discriminate between "old" target faces and "new" distactor faces. To begin, I first scored participants' "hits" and "false alarms". Hits are correct identifications of a previously seen target face. False alarms are incorrect identifications of new faces as old. I then calculated hit rates and false alarm rates by dividing the total number of hits and false alarms by the total number old and new faces respectively (Stanislaw & Todorov, 1999). I calculated *d'* by converting hit rates (H) and false alarm rates (FA) to z-scores, and subtracting the z-score for FA from the z-score for H (*d'* = z(H)-z(FA)) (Stanislaw & Todorov, 1999). In order to analyse whether participants demonstrated an ORB, I calculated *d'* separately for black and white faces for each participant. A higher *d'* indicates that a participant is better able to discriminate between targets and distractors i.e. they demonstrate higher recognition accuracy. Therefore, an ORB would be present when a participant demonstrates higher recognition accuracy for SR faces than CR faces.

Main analyses. All statistical analyses were conducted in SPSS (version 25) and Microsoft Excel. The significance level for all statistical analyses was set at $\alpha = .05$. All data was analysed at the 95% confidence interval. I downloaded the data for each survey separately from the Qualtrics website. I then scored the face recognition data and recorded it in an excel spreadsheet along with participants race and the condition they took part in. I also recorded participants' raw scores for the different self-report measures. I then calculated *d'* scores for black and white faces separately for each participant. I used the Iatgen web application to score the IAT data, and recorded these scores in Excel. For my main analysis of the ORB, I imported my data to SPSS and performed a 3X2X2 mixed-design ANOVA with condition (exclusion vs. inclusion vs. control) and race of participant (black vs. white) included as between–subjects factors. Race of face (*d'* for black vs. white faces) was analysed as the within-subjects factor. To analyse the relationship of the IAT with the ORB, I performed a 3X2 factorial ANCOVA with

ORB scores as the dependent variable, and condition and race of participant as the independent variable, with the IAT scores included as a covariate.

Ethical considerations

All participants had to read and respond to an informed consent form (Appendix E) before starting the study. If they did not agree to participate, the survey was ended. The consent form informed participants of their right to participate voluntarily, their right to withdraw and that all information would be confidential and anonymous. They were informed that there was a risk of discomfort or the possibility of negative feelings resulting from the procedures in the study. Participants were deceived, as they were led to believe that the *Cyberball* was an online game played with real players. This was done to better simulate the effects of real life exclusion. I included a detailed debriefing form (Appendix F) that fully outlined the procedures in the study and explained the rationale behind the use of the procedures. In the debriefing form, I explained the use of deception and informed participants that no real players were involved in the game aside from them. I also provided them with contact information if they had any complaints or questions, and informed them that if they needed any referral information I would provide it to them. Additionally, I included and a set of inclusion games at the end of the exclusion condition to minimise the impact of exclusion. Therefore, all participants in the experimental conditions took part in games where they were included. The Department of Psychology Ethics Committee granted ethical approval for the study.

Results

Manipulation check. For the exclusion group, participants' average scores were M = 5.44, SD = 1.55 and M = 5.39, SD = 1.55 for item 1 and item 2 respectively. For the inclusion condition, the average score was lower for both item 1 (M = 3, SD = 1.73) and item 2 (M = 3.05, SD 1.66). I performed two independent samples t-tests on these scores to determine if this difference was significant. There was a significant difference for scores on both item 1 and 2 between the two groups, indicating that the manipulation was successful, $p < .01$. Participants therefore felt more excluded and ignored in the exclusion condition than in the inclusion condition.

ORB analysis

All 92 participants completed the face recognition task. In the exclusion condition there were 11 black participants and 18 white participants. In the inclusion condition there were 14

black and 21 white, and in the control condition there were 8 black and 21 white participants. To determine if participants demonstrated biased recognition toward either white or black faces, I ran a repeated measures ANOVA with race of face (*d'* for white vs. black faces) as the within-subjects factor. I compared groups based on the participants' race to see if they showed a significant within-subject difference in their recognition of white and black faces. Prior to running any ANOVA, I checked P-P plots of *d'* scores to assess the assumption of normality (Appendix G). Both plots indicated that the data was approximately normally distributed. For black participants the within-subjects effect of race of face was significant, F (1,32) = 8.12, p = .008, *partial* n^2 = .2. However, for white participants the effect of race of face was not significant, F (1,58) = .33, p = .57, *partial* n^2 = .006. These results demonstrate that an ORB was indeed present, however it was only evident for black participants. Black participants demonstrated significantly better recognition of black faces (M = .49, SD = .67) than white faces (M = .13, SD = .56). White participants showed slightly better recognition of white faces (M = .4, SD = .51) compared to black faces (M = .33, SD = .71). Therefore, black participants showed a significant ORB, while white participants did not.

Main analysis

I ran a 3X2X2 mixed-design ANOVA to investigate the effects of exclusion and inclusion on the ORB. Condition (exclusion vs. inclusion vs. control) and race of participant (black vs. white) were included as between–subjects factors. Race of face (*d'* for black vs. white faces) was analysed as the within-subjects factor. Levene's test for equality of variance was non-significant for *d'* for black (p = .87) and white faces (p = .6), indicating the assumption of homogeneity of variance was upheld. Within-subjects effects indicated there was a significant main effect for race of face, F (1,86) = 4.27, p = .04, partial n^2 = .05. There was also a significant interaction effect of race of face by race of participant, F (1,86) = 7.53, p = .007, partial n^2 = .08. No other effects were significant. Pairwise comparisons for the main effect of race of face indicate that black faces (M = .46) were recognised significantly more accurately than white faces (M = .28) across conditions and racial groups. As shown earlier, the significant effect of the interaction of race of face by race of participant is a result of black participants overall better recognition of black faces than white faces.

Tests of between-subjects effects showed there was only a significant main effect for condition, F (2,86) = 12.47, p < .001, partial n^2 = .23. Pairwise comparisons for this effect

showed that participants' overall recognition accuracy in the exclusion condition (M = .21) was very similar to that of the inclusion condition (M = .18), p = 1. However, recognition performance in the control group (M = .7) was significantly higher than both the exclusion and inclusion group, p < .001. These results indicate that participants' performed much better in their overall recognition in the control group than both the inclusion and exclusion group. Participants' accuracy was the similar for the inclusion and exclusion condition, with marginally higher accuracy in the exclusion condition.

Table 1

Descriptive statistics of d' by condition and race of participant

	d' black			d' white		
	Exclusion	Inclusion	Control	Exclusion	Inclusion	Control
Black	.46 (.51)	.12 (.51)	1.12 (.6)	-.003 (.54)	.05 (.5)	.46 (.64)
White	.20 (.66)	.18 (.72)	.59 (.69)	.17 (.43)	.39 (.38)	.6 (.6)

Note. Standard deviations appear in parentheses below means

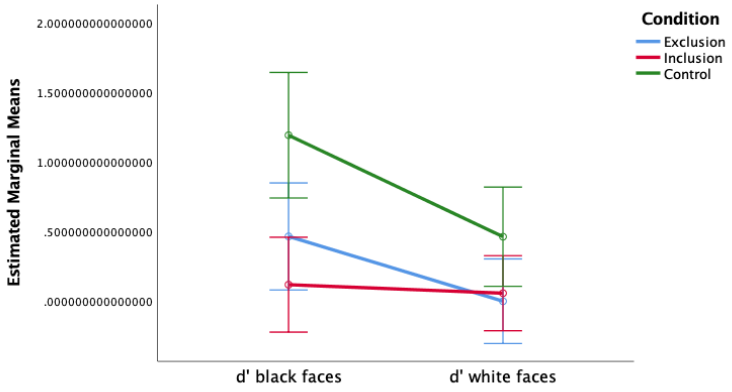

Figure 1. Estimated marginal means of *d'* for black and white faces by condition for black participants

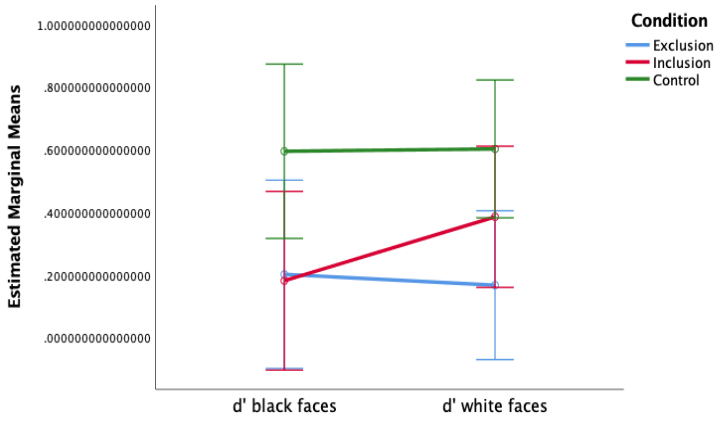

Figure 2. Estimated marginal means of *d'* for white and black faces by condition for white participants

Table 1 and figures 1 and 2 indicate that black participants in the exclusion condition showed better recognition of black faces. However, in the inclusion condition they showed similar recognition of black and white faces. In the control group black participants showed better recognition accuracy for black faces. White participants showed similar recognition for black and white players in the exclusion condition but better recognition of white players than black players in the inclusion condition. In the control condition, white participants showed almost equal recognition accuracy of black and white faces.

To test if any of these differences were significant I ran a repeated measures ANOVA with race of face as the within-subjects factor and compared participants by groups based on condition and race of participant. The only significant difference was for black participants' better recognition of black faces compared to white faces in the control condition, $p = .001$. In the exclusion condition black participants did demonstrate better recognition of black players, and although this result was not significant when compared to their recognition of white players, it is consistent with the hypothesis that exclusion increases the recognition of SR faces. While there was a significant ORB displayed by black participants in the control condition, only 8 participants were assigned to this condition, meaning the statistical power of this finding is low.

IAT analyses

IAT data was scored according to procedures outlined by Greenwald et al. (2003). Participants who responded too quickly (< 300ms) were dropped due to excessive speed. In total, 79 participants completed the IAT. I first ran a 2X3 factorial ANOVA to determine if there were

any differences in IAT scores across racial groups and conditions. Race of participant (black vs. white) and condition (exclusion vs. inclusion vs. control) were included as independent variables, while IAT scores were the dependent variable. Descriptive statistics for this analysis are shown in table 2. Levene's test of equality of variance was significant indicating the assumption of homogeneity of variance was upheld, $p = .3$. There was a main effect of race of participant, F (1,73) = 21.98, $p < .001$, partial $n^2 = .23$. No other effects were significant. Pairwise comparisons for this effect showed that white participants ($M = .41$) had significantly higher *d'* scores than black participants ($M = -.06$). These results indicate that white participants demonstrated a significantly greater association of white faces with positive attributes, indicating higher implicit prejudice. Black participants, however, demonstrated almost no implicit prejudice.

Table 2

Descriptive statistics of IAT scores by condition and race of participant

	IAT scores		
	Exclusion	Inclusion	Control
Black	-.12 (.31)	-.12 (.57)	.06 (-.27)
White	.48 (.33)	.35 (.35)	.4 (.5)
Total	.28 (.43)	.18 (.49)	.3 (.47)

Note. Standard deviations appear in parentheses below means

To test the relationship between IAT scores and the ORB, I calculated an ORB score for each participant by subtracting their *d'* for black faces from their *d'* for white faces (ORB = $d'_w - d'_b$) (Hills & Lewis, 2006). A positive number indicates there is an ORB, while a negative number indicates there is a bias toward CR faces (Hills, Cooper, & Pake, 2013). I then ran a 3X2 factorial ANCOVA, with condition (exclusion vs. inclusion v. control) and race of participant (black vs. white) as independent variables and ORB scores as the dependent variable. The effect of IAT was not significant, F (1,72) = .145, $p = .7$, partial $n^2 = .002$. Therefore, there was no significant difference between IAT scores and participants' ORB scores. This indicates that implicit prejudice does not appear to be associated with the ORB.

Self-report measures

A total of 79 participants completed the self-report measures. For all measures I averaged participants' responses to the different items to provide a mean response score for each measure for each participant (Finchilescu, 2010; Mackie et al., 2000).

Group identification. Black participants reported much higher group identification with black people (M = 6.24, SD =. 8) than white people (M = 3.36, SD = 1.25). White participants reported similar group identification with both white and black people, with slightly higher identification with white people (M = 4.79, SD = 1.07) than black people (M = 4.35, SD = .83). While there was no significant difference across conditions, black participants in-group identification was consistently high. This result seems consistent with the finding that black participants displayed better recognition of black faces.

Intergroup emotion. Black (M = 2.87, SD = 1.78) and white (M = 2.42, SD = 1.32) participants responded similarly, indicating both groups did not feel strong negative emotions toward out-group members.

Affective prejudice. Overall, both black and white participants displayed low affective prejudice. Black participants (M = 3.55, SD = 1.22) reported higher affective prejudice toward white people than white participants (M = 2.38, SD = .87) reported toward black people.

Correlation analysis

I ran a correlation analysis to analyse the relationships between all self-report measures, the IAT, and *d'* for white and black faces across participants. These correlations are shown in Table 3. These results indicate that out-group identification may be important in understanding levels of explicit prejudice toward out-group members. Similarly, intergroup emotion seems to be closely related to affective prejudice, as increases in negative emotions toward out-group members were associated with increased affective prejudice.

Table 3

Results of Pearson correlation analysis

Race of participant		Variables						
		1	2	3	4	5	6	7
Black	1. d' black faces	-	.325	.204	-.01	-.17	.01	-.32
	2. d' white faces		-	.26	.11	.23	-.27	-.15
	3. IAT scores			-	-06	-.001	-.07	.05
	4. Group identification with white people				-	-.13	-.65**	-.31
	5. Group identification with black people					-	.3	.13
	6. Affective prejudice						-	.5
	7.Intergroup emotion							-
White	1. d' black faces	-	.11	-.09	.11	.1	.03	-.04
	2. d' white faces		-	-.13	-.06	.17	.1	-.09
	3. IAT scores			-	.14	.1	.12	.2
	4. Group identification with white people				-	.47**	-.13	.11
	5. Group identification with black people					-	-.57**	-.38**
	6. Affective prejudice						-	.47**
	7.Intergroup emotion							-

Note. ** $p < .001$

Discussion

The hypothesis that participants would display an ORB was partially confirmed by the finding that black participants demonstrated an ORB. This finding differs from previous studies in that white participants in South African universities usually display an ORB (Seutloali, 2014; Wright et al., 2003). In previous studies of the ORB in South Africa, black participants have responded similarly for black and white faces (Seutloali, 2014; Wright et al., 2003). The small sample size for black participants limits the conclusions that can be drawn from these results, but a number of factors may explain this finding. Black people make up the majority of the population in South Africa, which means that black participants are likely to have increased contact and therefore greater perceptual expertise with black South African faces than with white faces (Wright et al., 2003). Additionally, while the lack of ORB for black participants was previously explained as possibly resulting from unequal social power, where white people have more social, economic, and political power, it may be that younger black South Africans have less contact with white people due to increasing social power for black South Africans (Shriver & Hugenberg, 2010). It may be that black South African have gained increased social power or representation in university spaces, meaning that there is a more proportionate representation of black people in spaces that were previously dominated by white people. The lack of ORB shown

by white participants is especially interesting and provides further support for this theory, as there was a larger sample of white participants. Changing social power may therefore mean that white participants minority status is resulting in more contact with black people and therefore increased expertise with black faces.

Limited support for the second hypothesis was provided by black participants' higher recognition of SR faces in the exclusion condition, although this finding was not significant. This recognition accuracy for SR players may provide some support for theories of negative contact and social exclusion, as black participants displayed higher mean recognition of black players than white players when white players excluded them. This finding is consistent with the theory that negative contact and social exclusion increased the salience of racial group membership and created feelings of in-group belonging, which may have led black participants to categorise white faces as out-group members (Barlow et al., 2012; Rutland & Killen, 2015). Therefore, following social cognitive theories, they would have processed white players categorically while individuating black players. In support of this, black participants identified more with black people than white people in general, displaying high in-group identification. When they were excluded by white players, increased category salience resulting from the experience negative contact (social exclusion) may have led black participants to identify more with other black players in the game, which in turn resulted in them individuating those players (Barlow et al., 2012; Rutland & Killen, 2015). However, another interpretation is that social exclusion reduced black participants SR recognition accuracy to the same level as their accuracy for CR players in the control group. This result may therefore rather indicate that the task was too difficult in the experimental conditions. The fact that black participants recognition of white faces was slightly negative in the inclusion condition suggests that the task was indeed too difficult (Stanislaw & Todorov, 1999).

Inclusion did not appear to have the same effect for black participants. It is possible that they did not take part in the exclusion of white players, instead passing the ball to them even though other black players did not. If this was case, category salience could have been reduced as participants included CR players. Providing some support for this theory, black participants showed almost no implicit prejudice, which is predictive of social exclusion of out-group members (Rutland & Killen, 2015). Conversely, white participants scored significantly higher in implicit prejudice and also demonstrated higher recognition of SR players in the inclusion

condition. It is therefore possible that white participants were more complicit in the exclusion of black players during the inclusion game, leading them to individuate white players and categorise black players as out-group members (Barlow et al., 2012; Rutland & Killen, 2015).

The final hypothesis that implicit prejudice was related to the ORB was not confirmed. There was no significant difference between IAT scores and ORB scores, indicating that IAT scores were not associated with the ORB. Therefore, implicit prejudice does not appear to be related to the ORB, a finding that is consistent with a number of other studies (Ferguson et al., 2001; Young et al., 2012). Additionally, white participants displayed significantly more implicit prejudice than black participant but did not display an ORB, while the opposite was true for black participants. This further points to the possibility that implicit prejudice is not associated with the ORB. Overall the study provided no support for the role of prejudice in the ORB, but did suggest that social exclusion and in-group identification and the interrelation between these two factors may play an important role in explaining the ORB.

Limitations and future directions

A major limitation for the study was the low sample size for black participants. Just over half as many black participants as white participants took part in the study. This low sample size significantly reduced statistical power, and limited the conclusions that can be drawn from the data. Another possible limitation was the use of composite pictures as the stimuli for the face recognition task. Using these faces meant I was unable to vary the pictures between encoding and recognition, which would have been especially important in ensuring the task was sufficiently difficult for the control group. Varying the colour (full colour to greyscale) and view of the face ensures participants are remembering the face rather than the picture (Gwinn, Barden, & Judd, 2015). Also related to the use of composite faces is the possibility that white faces were more difficult to recognise because they were more similar to each other. This may have occurred because the photo database used to generate the white faces was much smaller than the one used to generate black faces. The possibility that white faces were harder to distinguish could therefore have explained the lack of an ORB in white participants, and the presence of one for black participants. However, due to the nature of the study, it was important that no real people were implicated in race-based social exclusion, as there was a risk of participants recognising the person they believed had excluded them. Finally, most participants in the study completed it as an online survey, meaning that responses were less reliable as they were not

responding in a controlled environment. Additionally, screen sizes would have varied widely, thereby influencing the size of the images. Distributing the experiment as an online survey was important though, as it allowed for a larger number of responses to be collected in a short amount of time, significantly increasing the overall sample size.

Future studies should focus on the relationship between social exclusion and group identification in producing the ORB, as theories of group identification and negative contact appear to have received limited support from the findings of this study. Additionally, the absence of the ORB in white participants is especially interesting and future studies should attempt to replicate this finding, and explore the underlying factors that have produced this result. Future studies should also focus on ensuring that face sets used for recognition tasks are pilot tested to ensure the task is sufficiently difficult. This will prevent the possibility that one set of faces is more difficult to accurately discriminate between, and prevent the possibility that a recognition bias is created or diminished due to one set of faces being more similar than the other.

Summary and conclusion

The study provided limited support for the association between social exclusion and the ORB, but did not find any relationship between the ORB and prejudice. It is also interesting that SR recognition bias was found for black participants and not white participants, however these findings were limited by the small sample size. Research into the factors underlying the ORB has important implications in the criminal justice system in particular. Mistaken eyewitness identifications make up a significant number of false incarcerations, and any evidence of underlying causes of face recognition deficits may help to target the factors that lead to inaccurate identifications. Research into the factors underlying face recognition is therefore essential, as it has important implications for ensuring that mistaken identification does not lead to false imprisonments in the future.

Acknowledgments

I would like to thank my supervisor, Professor Colin Tredoux, for his help and guidance throughout the year.

I would also like to thank all the members of the eyewitness group for their feedback and help during meetings. I would like to thank Tania Wittwer for allowing me to use her Qualtrics account and showing me how to parts of set up the experiment in Qualtrics.

Lastly, I would like to thank everyone who participated in my study and shared my study with other people to help me get the responses I needed.

References

Barlow, F. K., Paolini, S., Pedersen, A., Hornsey, M. J., Radke, H. R. M., & Harwood, J.,…, (2012). The Contact Caveat: Negative Contact Predicts Increased Prejudice More Than Positive Contact Predicts Reduced Prejudice. *Personality and Social Psychology Bulletin*, *38*(12), 1629–1643. http://doi.org/10.1177/0146167212457953

Bernstein, M. J., Sacco, D., Young, S. G., & Hugenberg, K. (2014). The Impact of Race and Inclusionary Status on Memory for Ingroup and Outgroup Faces. *Basic and Applied Social Psychology*, *36*(3), 191–198. https://doi.org/10.1080/01973533.2014.887565

Brigham, J. C., & Barkowitz, P. (1978). Do "They all look alike?" The Effect of Race, Sex, Experience, and Attitudes on the Ability to Recognize Faces. *Journal of Applied Social Psychology*, *8*(4), 306–318. https://doi.org/10.1111/j.1559-1816.1978.tb00786.x

Carpenter, T., Pogacar, R., Pullig, C., Kouril, M., Aguilar, S., LaBouff, J. P., …, Chakroff, A. (2018). Conducting IAT Research within Online Surveys: A Procedure, Validation, and Open Source Tool. *PsyArXiv*. https://doi.org/10.31234/osf.io/hgy3z

Cohen, J. (1988). Statistical power analysis for the behavioral sciences. *Statistical Power Analysis for the Behavioral Sciences*. https://doi.org/10.1234/12345678

Ferguson, D. P., Rhodes, G., Lee, K., & Sriram, N. (2001). "They all look alike to me": Prejudice and cross-race face recognition. *British Journal of Psychology*, *92*(4), 567–577. http://doi.org/10.1348/000712601162347

Finchilescu, G. (2010). Intergroup anxiety in interracial interaction: The role of prejudice and metastereotypes. *Journal of Social Issues*, *66*(2), 334–351. https://doi.org/10.1111/j.1540-4560.2010.01648.x

Gaither, S. E., Pauker, K., Slepian, M. L., & Sommers, S. R. (2016). Social belonging motivates categorization of racially ambiguous faces. *Social cognition*, *34*(2), 97-118. https://doi.org/10.1521/soco.2016.34.2.97

Gawronski, B., & De Houwer, J. (2014). Implicit measures in social and personality psychology. *Social and Personality Psychology*, *1*(519), 1–28. https://doi.org/10.1017/CBO9780511996481.016

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: the implicit association test. *Journal of Personality and Social Psychology*, *74*(6), 1464–1480. https://doi.org/10.1037/0022-3514.74.6.1464

Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and Using the Implicit Association Test: I. An Improved Scoring Algorithm. *Journal of Personality and Social Psychology*, *85*(2), 197–216. https://doi.org/10.1037/0022-3514.85.2.197

Guinote, A., Willis, G. B., & Martellotta, C. (2010). Social power increases implicit prejudice. *Journal of Experimental Social Psychology*, *46*(2), 299–307. https://doi.org/10.1016/j.jesp.2009.11.012

Gwinn, J. D., Barden, J., & Judd, C. M. (2015). Face recognition in the presence of angry expressions: A target-race effect rather than a cross-race effect. *Journal of Experimental Social Psychology*, *58*, 1–10. https://doi.org/10.1016/j.jesp.2014.12.001

Hartgerink, C. H. J., Van Beest, I., Wicherts, J. M., & Williams, K. D. (2015). The ordinal effects of ostracism: A meta-analysis of 120 cyberball studies. *PLoS ONE*, *10*(5). https://doi.org/10.1371/journal.pone.0127002

Hills, P. J., Cooper, R. E., & Pake, J. M. (2013). Removing the own-race bias in face recognition by attentional shift using fixation crosses to diagnostic features: An eye-tracking study. *Visual Cognition*, *21*(7), 876–898. https://doi.org/10.1080/13506285.2013.834016

Hills, P. J., & Lewis, M. B. (2006). Reducing the own-race bias in face recognition by shifting attention. *Quarterly Journal of Experimental Psychology*, *59*(6). https://doi.org/10.1080/17470210600654750

Hugenberg, K., & Bodenhausen, G. V. (2003). Facing Prejudice: Implicit Prejudice and the Perception of Facial Threat. *Psychological Science*, *14*(6), 640–643. http://doi.org/10.1046/j.0956-7976.2003.psci_1478.x

Knowles, M. L., & Gardner, W. L. (2008). Benefits of membership: The activation and amplification of group identities in response to social rejection. *Personality and Social Psychology Bulletin*, *34*(9), 1200–1213. https://doi.org/10.1177/0146167208320062

Lebrecht, S., Pierce, L. J., Tarr, M. J., & Tanaka, J. W. (2009). Perceptual other-race training reduces implicit racial bias. *PLoS ONE*, *4*(1). http://doi.org/10.1371/journal.pone.0004215

Ma, Y., Yang, S., & Han, S. (2011). Attitudes influence implicit racial face categorization in a perceptual task. *Group Processes and Intergroup Relations*, *14*(6), 887–899. http://doi.org/10.1177/1368430211409962

Mackie, D. M., Devos, T., & Smith, E. R. (2000). Intergroup emotions: Explaining offensive action tendencies in an intergroup context. *Journal of Personality and Social Psychology*, *79*(4), 602–616. https://doi.org/10.1037/0022-3514.79.4.602

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*. http://doi.org/10.1037//1076-8971.7.1.3

McGugin, R. W., Tanaka, J. W., Lebrecht, S., Tarr, M. J., & Gauthier, I. (2011). Race-Specific Perceptual Discrimination Improvement Following Short Individuation Training With Faces. *Cognitive Science*, *35*(2), 330–347. https://doi.org/10.1111/j.1551-6709.2010.01148.x

Qian, M. K., Quinn, P. C., Heyman, G. D., Pascalis, O., Fu, G., & Lee, K. (2017). Perceptual individuation training (but Not Mere Exposure) reduces implicit racial bias in preschool children. *Developmental Psychology*, *53*(5), 845–859. http://doi.org/10.1037/dev0000290

Paolini, S., Harwood, J., & Rubin, M. (2010). Negative intergroup contact makes group memberships salient: Explaining why intergroup conflict endures. *Personality and Social Psychology Bulletin*, *36*(12), 1723–1738. https://doi.org/10.1177/0146167210388667

Rutland, A., & Killen, M. (2015). A developmental science approach to reducing prejudice and social exclusion: Intergroup processes, social-cognitive development, and moral reasoning. *Social Issues and Policy Review*, *9*(1), 121–154. https://doi.org/10.1111/sipr.12012

Sacco, D. F., Wirth, J. H., Hugenberg, K., Chen, Z., & Williams, K. D. (2011). The world in black and white: Ostracism enhances the categorical perception of social information. *Journal of Experimental Social Psychology*, *47*(4), 836–842. https://doi.org/10.1016/j.jesp.2011.03.001

Scroggins, W. A., Mackie, D. M., Allen, T. J., & Sherman, J. W. (2016). Reducing Prejudice With Labels: Shared Group Memberships Attenuate Implicit Bias and Expand Implicit Group Boundaries. *Personality and Social Psychology Bulletin*, *42*(2), 219–229. https://doi.org/10.1177/0146167215621048

Seutloali, G. M. (2014). Own-race bias in facial recognition amongst black, coloured and white participants. *Unpublished honour's thesis). University of Cape Town, South Africa*.

Shriver, E. R., & Hugenberg, K. (2010). Power, individuation, and the cross-race recognition deficit. *Journal of Experimental Social Psychology*, *46*(5), 767–774. http://doi.org/10.1016/j.jesp.2010.03.014

Shriver, E. R., Young, S. G., Hugenberg, K., Bernstein, M. J., & Lanter, J. R. (2008). Class, race, and the face: Social context modulates the cross-race effect in face recognition. *Personality and Social Psychology Bulletin*, *34*(2), 260–274. http://doi.org/10.1177/0146167207310455

Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments, and Computers*, *31*(1), 137–149. https://doi.org/10.3758/BF03207704

Teige-Mocigemba, S., Klauer, K. C., & Sherman, J. W. (2010). A practical guide to Implicit Association Tests and related tasks. In *Handbook of Research Methods in Social and Personality Psychology* (pp. 117–139). https://doi.org/10.1016/S1572-1000(05)00006-2

Tredoux, C., & Finchilescu, G. (2010). Mediators of the contact-prejudice relation among South African students on four university campuses. *Journal of Social Issues*, *66*(2), 289–308. https://doi.org/10.1111/j.1540-4560.2010.01646.x

Van Harmelen, A. L., Hauber, K., Moor, B. G., Spinhoven, P., Boon, A. E., Crone, E. A., & Elzinga, B. M. (2014). Childhood emotional maltreatment severity is associated with dorsal medial prefrontal cortex responsivity to social exclusion in young adults. *PLoS ONE*, *9*(1). https://doi.org/10.1371/journal.pone.0085107

Walker, P. M., & Hewstone, M. (2008). The influence of social factors and implicit racial bias on a generalized own-race effect. *Applied Cognitive Psychology*, *22*(4), 441–453. http://doi.org/10.1002/acp.1382

Williams, K. D., Cheung, C. K. T., & Choi, W. (2000). Cyberostracism: Effects of being ignored over the Internet. *Journal of Personality and Social Psychology*, *79*(5), 748–762. https://doi.org/10.1037/0022-3514.79.5.748

Williams, K. D., Yeager, D. S., Cheung, C. K. T., & Choi, W. (2012). Cyberball (version 4.0) [Software]. Available from https://cyberball.wikispaces.com.

Wright, D. B., Boyd, C. E., & Tredoux, C. G. (2003). Inter-racial contact and the own-race bias for face recognition in South Africa and England. *Applied Cognitive Psychology*, *17*(3), 365–373. https://doi.org/10.1002/acp.898

Young, S. G., & Hugenberg, K. (2012). Individuation motivation and face experience can operate jointly to produce the own-race bias. *Social Psychological and Personality Science*. http://doi.org/10.1177/1948550611409759

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*. http://doi.org/10.1177/1088868311418987

Zanna, M. P. (1994). On the nature of prejudice. *Canadian Psychology/Psychologie Canadienne*, *35*(1), 11–23. https://doi.org/10.1037/0708-5591.35.1.11

Appendix A

In-game screenshot of *Cyberball*

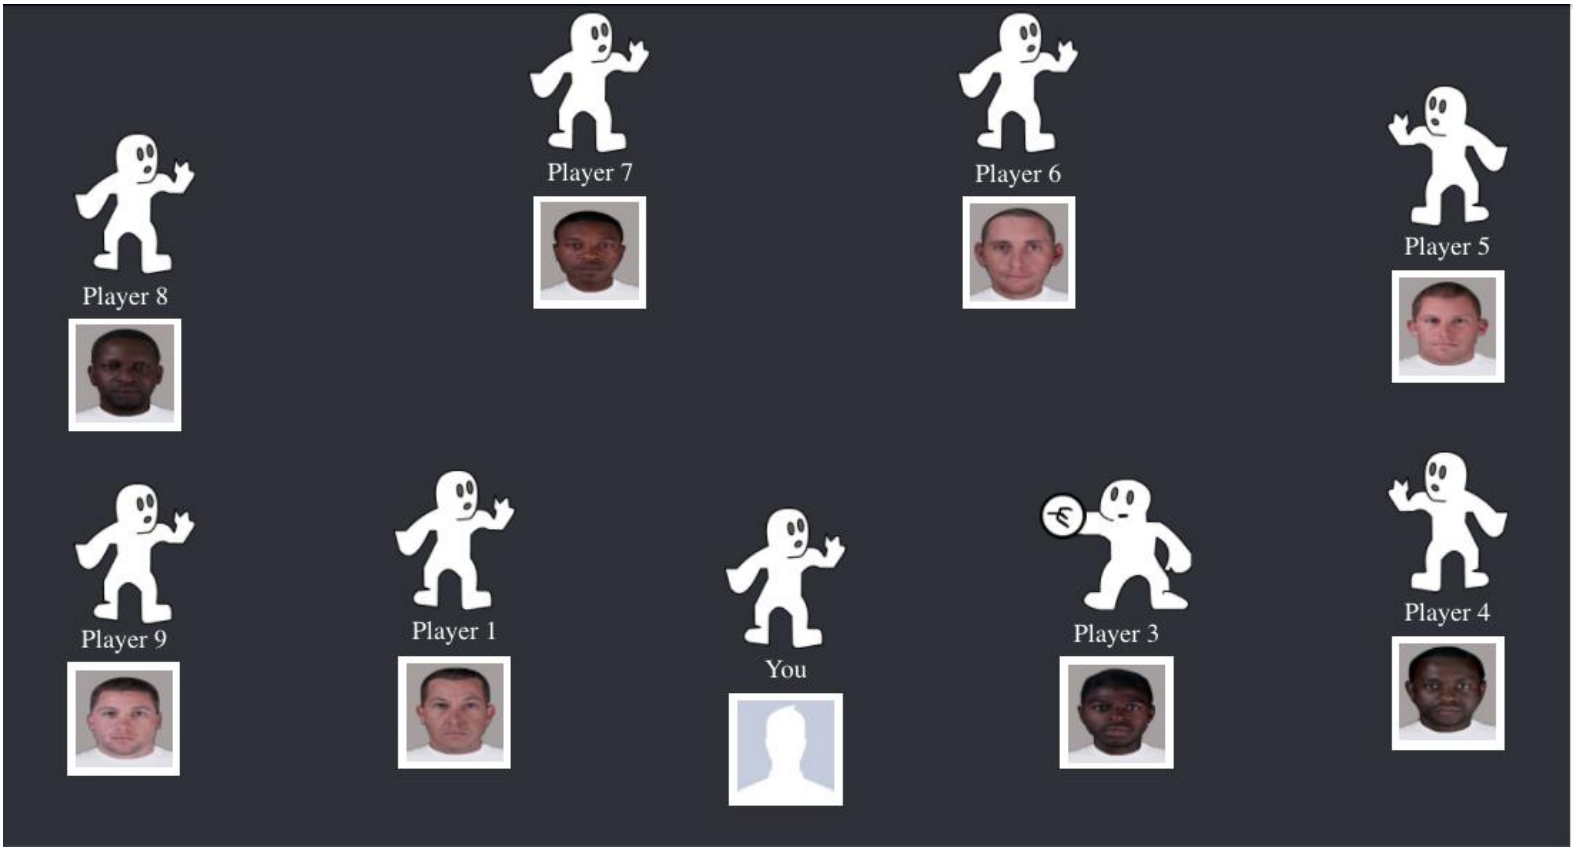

Appendix B

Group identification scale

Please respond to the following four statements with regard to white people (select the response that best reflects your agreement with the statement):

	Strongly disagree	Disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Agree	Strongly agree
1. You stand by them	○	○	○	○	○	○	○
2. You share their values	○	○	○	○	○	○	○
3. You feel similar to them	○	○	○	○	○	○	○
4. You feel close to them	○	○	○	○	○	○	○

Please respond to the following four statements with regard to black people (select the response that best reflects your agreement with the statement):

	1. Strongly disagree	Disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Agree	7. Strongly agree
5. You stand by them.	○	○	○	○	○	○	○
6. You share their values	○	○	○	○	○	○	○
7. You feel similar to them	○	○	○	○	○	○	○
8. You feel close to them	○	○	○	○	○	○	○

Appendix C

Affective prejudice scale

Please select the number, for each pair of adjectives, that best applies to your responses to the following statement:

In general, I feel the following ways, emotionally, towards white people:

1. Warm	○ ○ ○ ○ ○ ○ ○	7. Cold
1. Positive	○ ○ ○ ○ ○ ○ ○	7. Negative
1. Friendly	○ ○ ○ ○ ○ ○ ○	7. Hostile
1. Trusting	○ ○ ○ ○ ○ ○ ○	7. Suspicious
1. Respect	○ ○ ○ ○ ○ ○ ○	7. Contempt
1. Admiration	○ ○ ○ ○ ○ ○ ○	7. Disgust

Appendix D

Intergroup emotion scale

Select the number that best describes your agreement the following statements:

Please indicate the degree to which you feel the following emotions toward white people:

	Strongly disagree	Disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Agree	Strongly agree
1. I feel angry	○	○	○	○	○	○	○
2. I feel displeased	○	○	○	○	○	○	○
3. I feel irritated	○	○	○	○	○	○	○
4. I feel furious	○	○	○	○	○	○	○
5. I feel worried	○	○	○	○	○	○	○
6. I feel anxious	○	○	○	○	○	○	○
7. I feel afraid	○	○	○	○	○	○	○
8. I feel fearful	○	○	○	○	○	○	○

Appendix E

Informed consent form

Overview
I am a Psychology Honours student at the University of Cape Town investigating the effects of an online computer game on participants' performance in a few short tasks. In the experiment, participants assigned to the experimental condition will play a simple online computer game with other players. In the game, you will throw a ball virtual ball around with the other players for a few minutes. Following this, you will complete a brief task and answer a few short questionnaires.

Procedure
If you decide to participate, you will be randomly assigned to one of three conditions. All procedures will be completed in this survey. The study will take about 30 to 45 minutes to complete.

Risks
If you are assigned to one of the experimental groups, there are potential negative emotional effects related to the procedures in the study. You are free to withdraw at any point that you feel uncomfortable or unable to continue. You can also refuse participation at any point in the study.

Benefits
There are no direct benefits to you as a participant in the study. However, this research has a number of important practical implications and can be used to further our understanding of different psychological processes.

Voluntary Participation and Right to Withdraw
Participation in the study is completely voluntary and you will face no penalties at all for not participating in the study. If you decide to participate, you are free to withdraw from the study at any point without fear of any negative consequences. You will not have to provide any reasons for your decision to withdraw from the study.

Confidentiality
All information gathered about you in the study will be completely confidential, and your name will not appear anywhere in the study. All data will be kept on a password-protected computer. Reports about the study will not identify any of the participants in the study.

Exclusion Criteria
Due to the nature of the study and feelings of discomfort, sadness or anxiety that may arise, you will not be able to participate if you have a current or previous diagnosis of a psychological disorder, such as depression or anxiety related disorders (or any disorder covered in the DSM).

Questions
If you have any further questions about the study or would like any additional information, please contact me:

Daniel Derbyshire
Email: DRBDAN001@myuct.ac.za

Please contact Rosalind Adams (rosalind.adams@uct.ac.za or 021 650 3417) if you have any complaints or comments about the study or me as a researcher.

Appendix F

Debriefing form

Investigating the effect of prejudice on the cross-race effect in face recognition

Dear participant,

Thank you for your participation in the study! This form will provide an outline of the study and explain the procedures involved in the study.

1. Purpose of the study

The aim of the study was to investigate racial differences in the perception and memory of faces and whether or not this is related to participants' implicit prejudice. The cross-race effect refers to people's ability to better remember or recognise faces that are part of the same racial group (same-race faces) as theirs compared to faces of a different race (cross-race faces). This effect has been found to reliably occur in many different contexts. There are a number of different theories about the underlying reasons for this effect, with perceivers' implicit racial prejudice representing one possible factor that may explain differences in people's ability to recognise same-race versus cross-race faces. Usually facial recognition tasks are done in laboratory environments and involve passive viewing of faces of different races. Therefore, a further aim of the study was to investigate the relationship between prejudice and face recognition when participants were made actively aware of different racial groups during the study. For this reason, if you were assigned to one of the experimental conditions you were either included by players of the same racial group as you or excluded by those of a different race during the Cyberball games. Additionally, in the exclusion condition, cross-race players also excluded players of the same race as you, while in the inclusion condition same-race players included you but cross-race players were excluded. This active awareness of racial groups or differences between racial groups was intended to facilitate belonging or identification with your own racial in-group in order to assess how this would affect your memory of same-race and cross-race faces. In the control condition, participants simply performed a facial recognition task followed by the other measures. Through studying the cross-race effect in a way that has participants' actively identifying with their own racial group, it is hoped that a better understanding of the cross-race effect, and the role of prejudice and social factors, can be gained.

2. Deception

In the experimental conditions, participants played a game called Cyberball, which is commonly used to experimentally manipulate social exclusion and inclusion. In this game, you were either excluded or included by members of the same race or different race. The other players in the game were not real players but are in fact programmed, computer players. The deception was used to better simulate real beliefs and feelings of exclusion, but no real players were involved in the game except you as a participant.

3. Negative effects associated with exclusion

All participants completed an inclusion condition to attempt to control for the negative emotional effects related to the experience of exclusion. However, exclusion can induce negative feelings

and emotions and if you feel that you need any additional support or information following the exclusion condition, please let me know and I will provide you with any additional information or referral information that you may need.

4. Further requirements of you as a participant

Please try not to share any information about the study that you learned through participating in the study with other students as this can bias their responses if they decide to participate in the study.

5. Reminder of confidentiality

I would just like to remind you that all your information will be kept confidential and stored in a locked cabinet. Your name only appears on the consent form and will not be linked to any of the data collected about you.

Thank you once again for your participation in the study! If you have any further questions please email me, Daniel Derbyshire, on DRBDAN001@myuct.ac.za.

If you have any complaints or comments about the study or about the way you were treated during the study, please contact Rosalind Adams on rosalind.adams@uct.ac.za or 021 650 3417.

Appendix G

PP-plots of *d'* for black and white faces

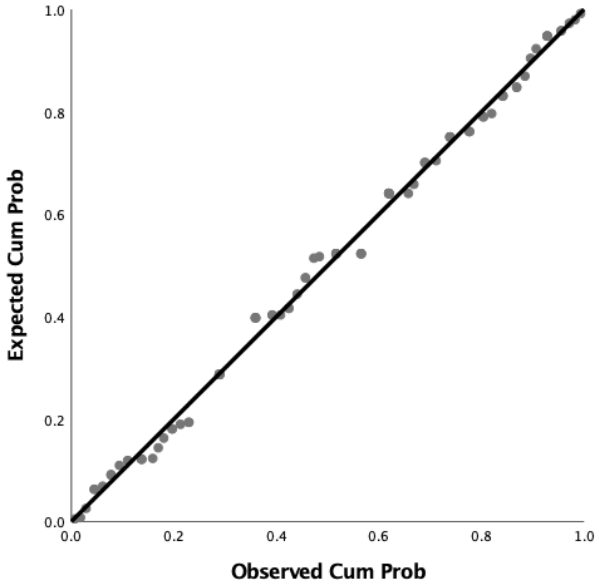

Figure 1. Normal PP plot of *d'* for black faces

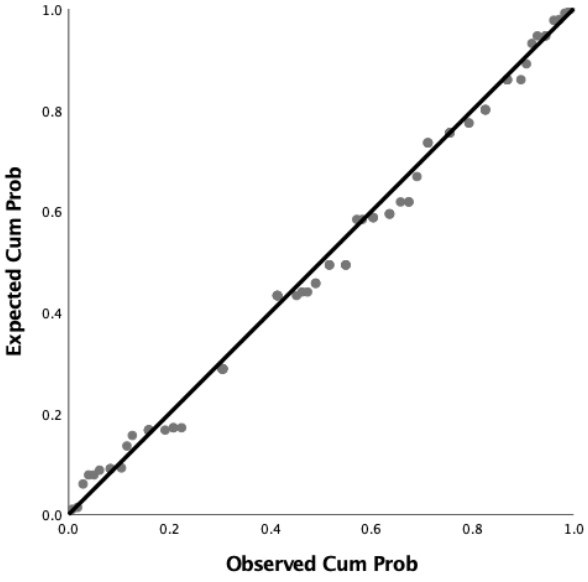

Figure 2. Normal P-P plot of *d'* for white faces

Appendix H

Ethics approval letter

UNIVERSITY OF CAPE TOWN

Department of Psychology

University of Cape Town Rondebosch 7701 South Africa
Telephone (021) 650 3417
Fax No. (021) 650 4104

03 August 2018

Daniel Derbyshire
Department of Psychology
University of Cape Town
Rondebosch 7701

Dear Daniel

I am pleased to inform you that ethical clearance has been given by an Ethics Review Committee of the Faculty of Humanities for your study, Prejudice and the own-race bias in face recognition. The reference number is PSY2018 -056.

I wish you all the best for your study.

Yours sincerely

University of Cape Town
PSYCHOLOGY DEPARTMENT
Upper Campus
Rondebosch

Lauren Wild (PhD)
Associate Professor
Chair: Ethics Review Committee