Out-Group Flies in the In-Group's Ointment:
Evidence of the Motivational Underpinnings of the In-Group Overexclusion Effect

Mark Rubin
and
Stefania Paolini
The University of Newcastle, Australia

This self-archived version is provided for non-commercial and scholarly purposes only.

The APA (6th ed) style reference for this article is as follows:

Rubin, M., & Paolini, S. (2014). Out-group flies in the in-group's ointment: Evidence of the motivational underpinnings of the in-group overexclusion effect. *Social Psychology*, 45, 265-273. doi: [10.1027/1864-9335/a000171](https://doi.org/10.1027/1864-9335/a000171)

We are grateful to the following people for their assistance in conducting this research:
Beatrice Bora, Mary-Claire Hanlon, Dane Poboka, and Amy Richards.

Correspondence concerning this article should be addressed to Mark Rubin at the School of Psychology, the University of Newcastle, Callaghan, NSW 2308, Australia. Tel: +61 (0)2 4921 6706. Fax: +61 (0)2 4921 6980. E-mail: Mark.Rubin@newcastle.edu.au

Abstract

People tend to misclassify ambiguous individuals as members of the out-group rather than the in-group. This *in-group overexclusion effect* (IO effect) is thought to occur because people are motivated to maintain their in-group's positivity by protecting it from potential out-group intrusions. The present research tested this explanation by asking university students ($N = 122$) to complete a self-esteem scale and then recall the group memberships of individuals who belonged to minimal groups. Consistent with predictions, participants misassigned significantly fewer individuals to the in-group than to the out-group when the in-group was positive and the out-group was negative but not when these valences were reversed. In addition, self-esteem negatively predicted the IO effect. Alternative explanations of the IO effect are discussed.

KEYWORDS: in-group overexclusion effect, social identity, self-esteem, minimal groups

Out-Group Flies in the In-Group's
Ointment:
Evidence of the Motivational Underpinnings
of the In-Group Overexclusion Effect

The *in-group overexclusion effect* (IO effect; Leyens & Yzerbyt, 1992) is an intergroup effect in which people have more stringent criteria and less proclivity for classifying ambiguous individuals as members of the in-group rather than the out-group. So, for example, people ask for a relatively large amount of information before classifying individuals as in-group members and less information before classifying individuals as out-group members (Dazzi, Voci, Brambilla, & Capozza, 1996; Leyens & Yzerbyt, 1992). People are also more likely to misclassify in-group members as out-group members rather than vice versa (e.g., Yzerbyt, Leyens, & Bellour, 1995). In particular, the IO effect has been observed when people make judgements of group membership (e.g., Northern Italian vs. Southern Italian) based on people's faces, some of which have been digitally morphed to contain a mixture of in-group and out-group stereotypical features (Boccatto, Capozza, & Falvo, 2003; Capozza, Boccatto, Andrighetto, & Falvo, 2009; Castano, 2004; Castano, Yzerbyt, Bourguignon & Seron, 2002; Knowles, & Peng, 2005; Pauker, Weisbuch, Ambady, Sommers, Adams, & Ivcevic, 2009).

Leyens and Yzerbyt (1992) proposed a motivational explanation of the IO effect. Following social identity theory (Tajfel & Turner, 1979), they assumed that people are motivated to create, maintain, and protect a positive social identity by creating, maintaining, and protecting the positivity of the in-group relative to comparable out-groups. Based on this assumption, they proposed that the IO effect occurs because people are motivated to protect their in-group from intrusion (contamination) by

negatively-valued out-group members. In other words and by analogy, people are motivated to keep out-group flies out of their in-group ointment.

Previous Tests of the Motivational Explanation of the IO Effect

Several studies have tested Leyens and Yzerbyt's (1992) motivational explanation of the IO effect. However, the results of these studies have been inconclusive.

In two separate studies, researchers at the University of Padova experimentally manipulated the need to protect social identity (Dazzi et al., 1996) and threat to in-group distinctiveness (Boccatto et al., 2003). Dazzi et al. (1996) manipulated the need to protect social identity by asking some participants to select a leader for the in-group and other participants to select a leader for the out-group. Although they replicated Leyens and Yzerbyt's (1992) results, their experimental manipulation had no significant effect. Boccatto, Capozza, and Falvo (2003) experimentally manipulated a threat to in-group distinctiveness by presenting information about the extent to which the in-group and out-group were similar to one another in terms of a range of behaviours. Again, this manipulation had no significant effect on the IO effect. Hence, these studies do not provide any supportive evidence for the motivational explanation.

Consistent with the motivational explanation, Castelli, Gorrasi, and Arcuri (2000) found a significant IO effect in a study in which the intergroup distinction was relevant to participants and a nonsignificant effect in a study in which the distinction was of low importance to participants. Critically, however, these researchers did not experimentally manipulate relevance within the same study and, consequently, the reliability of this moderating effect is unclear.

The most convincing evidence for the motivational account comes from a study by Castano et al. (2002). These researchers measured Northern Italian participants' in-group identification with their region (e.g., "I identify with Northern Italians"). Participants then categorised a series of faces as either Northern or Southern Italian. The researchers found that high identifiers were more likely than low identifiers to categorize the faces as Southern Italian. In other words, consistent with the social identity explanation, in-group identification moderated the size of the IO effect. This moderating effect has also been demonstrated in relation to the categorization of Black and White faces (Knowles & Peng, 2005, Study 3).

This previous research suggests that social identity is implicated in the IO effect, however, it also leaves some important questions unanswered. Although we now know that people are more likely to show the IO effect when they identify with their in-group, it remains unclear *why* they show this effect. The assumption is that people attempt to protect the in-group's positivity and, consequently, meet their need for self-esteem (Leyens & Yzerbyt, 1992). However, this assumption has never been tested directly. It is important to carry out this type of test in order to rule out other plausible explanations of the IO effect that also predict a relation with in-group identification. In particular, it is possible that the IO effect occurs because people are motivated to limit the size of the in-group in order to secure sufficient material resources for in-group members (Sherif, 1967) or achieve a more distinctive in-group (Livingstone, Spears, Manstead, & Bruder, 2011). It is also possible that the effect occurs because people are more familiar with in-group exemplars than they are with out-group exemplars (Corneille, Hugenberg, & Potter, 2007). All of these explanations

would predict that the IO effect is moderated by in-group identification because high identifiers are more concerned about their group's resources and distinctiveness and more attentive to their group's exemplars. Hence, the fact that in-group identification moderates the IO effect does not necessarily imply exclusive support for Leyens and Yzerbyt's (1992) motivational explanation.

Direct Tests of the Motivational Explanation

In the present research, we aimed to provide a more direct test of the motivational explanation. We predicted that if the IO effect is caused by the need to protect the in-group's positivity, then the effect would be strongest when the in-group has a positive valence and the out-group has a negative valence because, in this situation, group members would be motivated to keep negative out-group members out of their positive in-group. However, the IO effect should be weakened, nullified, or even reversed when the in-group has a negative valence and the out-group has a positive valence because, in this situation, group members would not be motivated to keep positive out-group members out of their negative in-group. This prediction is similar to Yzerbyt, Castano, Leyens, and Paladino's (2000) idea that the IO effect will be larger for high status groups than for low status groups.

We also tested a second prediction based on the social identity motivational explanation: If the IO effect is caused by the need for self-esteem, then people who have low self-esteem would be most likely to display the effect because they have the greatest need for self-esteem. Similar reasoning has been used to predict that people with low self-esteem will show the greatest prejudice and discrimination (e.g., Abrams & Hogg, 1988). Notably, however, the evidence for this self-esteem hypothesis

has been equivocal (Rubin & Hewstone, 1998). One reason for this limited evidence is that prejudice and discrimination represent relatively direct and blatant forms of self-enhancement, and people with low self-esteem prefer more indirect and subtle forms of self-enhancement because they lack the confidence to engage in more direct forms (Brown, 1988). In the present research, we assumed that the relation between self-esteem and the IO effect might be more reliable than that between self-esteem and discrimination because the IO effect represents a more subtle and indirect form of in-group bias that does not involve explicit evaluative judgements of people or groups.

Overview of the Present Research

In summary, we predicted that the IO effect would be strongest when the in-group has a positive valence and the out-group has a negative valence and weakened, nullified, or reversed when the in-group has a negative valence and the out-group has a positive valence. In addition, we predicted that people with low self-esteem would be more likely to display the IO effect than people with high self-esteem.

To test our first prediction, we varied the valence of the in-group and out-group (positive/negative) in a within-participants design and measured participants' recall of which target individuals belonged to which group. We expected that participants would misassign significantly more in-group targets to the out-group than vice versa when the in-group was positive and the out-group was negative but not when the in-group was negative and the out-group was positive. In other words, we expected the IO effect to occur only when it benefitted participants' social identity by excluding potentially negative out-group members from a positive in-group.

Note that we deliberately manipulated in-group valence and out-group

valence in a nonorthogonal manner: The in-group always had a positive valence when the out-group had a negative valence and vice versa. This approach of manipulating *intergroup* differences in valence was more appropriate than an orthogonal manipulation in which in-group valence varied independent of out-group valence. In particular, our approach was consistent with social identity theory, which assumes that groups establish their status and valence via intergroup comparisons with out-groups rather than via noncomparative assessments.

To test our second prediction, we measured participants' self-esteem as a potential predictor of the IO effect. We expected to find a negative relation between self-esteem and the IO effect: The lower people's self-esteem, the greater their need for self-esteem, and so the greater their motive to protect their in-group identity by excluding potential out-group members from the in-group.

It is important to note that both of the above predictions are specific to Leyens and Yzerbyt's (1992) motivational explanation, and neither can be derived from explanations based on material resources, the need for distinctiveness per se, or differential intergroup familiarity. Nonetheless, in order to completely rule out these alternative explanations from our research, we used minimal groups that had no connection with material resources, fixed and equal sizes, and equal familiarity due to their novelty and anonymity of group members.

One drawback of using the minimal group paradigm is that it was unsuitable for implementing the standard classification task that has been used to investigate the IO effect in previous research. A defining feature of the minimal group paradigm is that it does not include a stereotypical association between group members' personal information – usually code

numbers (e.g., Person number 12) – and their group membership (e.g., Group A or Group B). Given that classification tasks rely on this stereotypical association in order to operate, we were unable to use a standard classification paradigm to test our predictions. So, for example, it is reasonable to ask participants to classify individuals as Northern Italian or Southern Italian based on their facial features because Northern Italians are stereotyped as having fairer skin than Southern Italians. However, it is not reasonable to ask participants to classify individuals as members of Group A or Group B based on the individuals' code numbers (e.g., Person 3, Person 12) because these code numbers are not stereotypical of groups (e.g., people with higher numbers should not tend to belong to one group more than another). It is this lack of stereotypical association with personal information that helps to make minimal groups "minimal".

Given this issue, we used an alternative method of testing the IO effect that is based on the recall of the group memberships of briefly presented minimal group members. In this novel paradigm, we presented participants with information that indicated which of 20 code numbers belonged to which of two minimal groups. After a brief period, we removed this information and then asked participants to recall the group to which a particular code number had belonged. In this situation, the IO effect would be evident if participants incorrectly recalled more in-group members as belonging to the out-group than vice versa. We presented participants with numerous trials such as this. However, we varied the assignment of code numbers to groups during each presentation in order to preclude a stereotypical association between code numbers and groups. This approach allowed us to operationalize the IO effect in a minimal group paradigm and, consequently, to rule out alternative

explanations that may explain this effect in real groups.

Method

Participants

Participants were 122 domestic students at a large Australian university. The sample contained 48 men and 74 women. Participants had a mean age of 23.57 years ($SD = 5.14$) and ranged from 17 to 45 years.

Procedure and Measures

The research project was titled "memory task in social groups". Participants were told that the research was investigating implicit learning and memory recall for people in social groups who are placed in positive and negative contexts.

Participants attended individual research sessions. They began the session by drawing a card out of a bag at random. The card assigned them a number ostensibly from 1 to 20 but actually either 3 or 14. Participants retained the same number throughout the study (3 or 14) and used this number to identify themselves during the study.

Participants then completed a measure of self-esteem. Following Brown's (1988) investigation of indirect self-enhancement, we asked participants to complete a measure of global personal self-esteem in the form of Rosenberg's (1965) 10-item Self-Esteem Scale. An example item is "I feel that I have a number of good qualities."

The main part of the study consisted of a memory recall task. During each trial of this task, a computer screen presented a diagram that was similar to Figure 1.

Figure 1. Example of the Research Stimuli Presented as Part of the Memory Recall Task.

The diagram showed two groups of people, with each person represented by a code number from 1 to 20. One group was located on the left side of the computer screen, and the other group was located on the right side of the screen. There were ten people in each group. The participants' code number (3 or 14) was included in one of the groups. Hence, one of the groups represented an in-group, and the other represented an out-group.

To manipulate the valence of each group, we designed the left side of the screen to look like a bucket that contained soapy water and the right side of the screen to look like a dustbin that contained rubbish. These objects were chosen in order to capture participants' attention and engage them with the valence factor. To reinforce the meaning of this valence factor, the bucket was labelled "clean" and the dustbin was labelled "dirty". We also instructed participants to consider people in the bucket as clean and people in the dustbin as dirty.

During each trial presentation, participants were given 5 seconds to memorize which people belonged to which group. The diagram was then removed, a person's code number was presented, and participants were asked to recall which group the person had belonged to.

Participants responded by pressing the "A" key on the extreme left of their computer keyboard if they believed that the target person belonged to the clean group and the "L" key on the extreme right of their keyboard if they believed that the target person belonged to the dirty group. In order to increase the association between these keys and the clean and dirty groups, labels were affixed to the top of each key, with the "A" key being labelled "C" for "clean" and the "L" key being labelled "D" for "dirty".

Participants completed 2 practice trials, 64 experimental trials, and 8 self-assignment trials. In each trial, different code numbers were assigned to different groups. Hence, the people in each participant's in-group and the associated out-group changed from trial to trial, making it impossible for participants to predict who would be in their group during each trial. Participants were told that this assignment process was random. In fact, 10 different code numbers were always assigned to the clean group and 10 to the dirty group. These equal group sizes ruled out the influence of in-group distinctiveness effects.

Participants' own code numbers were assigned to the clean and dirty groups an equal number of times during the 64 experimental trials. This part of the procedure manipulated the valence of the in-group and out-group. Hence, on some trials the participant's in-group was positive (clean bucket) and the out-group was negative (dirty dustbin), and on other trials the in-group was negative (dirty dustbin) and the out-group was positive (clean bucket).

Different code numbers were selected as targets-to-be-recalled in different trials. Specifically, participants were asked to recall the location of 16 targets who were in the clean bucket with them (positive in-group targets), 16 who were in the dirty dustbin with them (negative in-group

targets), 16 who were in clean bucket when they were in the dirty dustbin (positive out-group targets), and 16 who were in the dirty dustbin when they were in the clean bucket (negative out-group targets). Hence, we used 64 trials to test the effects of a 2 (target group: in-group/out-group) $\times$ 2 (group valence: positive/negative) within-subjects design on participants' memory recall.

The 64 targets-to-be-recalled were presented in a single fixed random order in order to give participants the impression of a random selection. To further support this impression, we included eight additional self-assignment trials in which participants were asked to recall the location of their own code number. As with the experimental trials, participants pressed the "A" key (labelled "C") if they believed that their code number belonged to the clean bucket group and the "L" key (labelled "D") if they believed that it belonged to the dirty dustbin group. These self-assignment trials were interspersed randomly among the other trials.¹

Feedback from participants during a pilot study revealed that some participants used the strategy of memorizing the targets in only one of the two groups. They then combined memory recall with the process of elimination in order to improve their recall accuracy. In order to reduce the error variance associated with between- and within-subject variation in the use of this strategy, we instructed participants to spend an equal amount of time looking at the code numbers in both groups. A subsequent pilot study showed that these revised instructions produced a 30% error rate. Hence, our task made it difficult for participants to recall who belonged to which group.

After completing the memory recall trials, we asked participants to estimate how many times their own code number had appeared in the clean and dirty groups on a scale from 0 to 72. We used this

retrospective measure of group membership to establish how accurate participants were in recalling the location of the self in each group.

Finally, participants completed Rubin, Paolini, and Crisp's (2010) Perceived Awareness of the Research Hypothesis scale. This 4-item scale measures the extent to which participants believe that they are aware of the research hypothesis, and it is used to establish the degree to which demand characteristics influence research results. Example items are "I knew what the researchers were investigating in this research," and "I wasn't sure what the researchers were trying to demonstrate in this research" (reverse scored).

Results

Retrospective Recall of Group Membership

Participants were asked to recall the number of times that their own code number appeared in the clean and dirty groups across all of the trials (experimental and self-assignment). The self appeared in the clean bucket 40 times (8 times in the self-assignment trials and 32 times in the experimental trials) and 32 times in the dirty bucket (32 times in the experimental trials). In order to obtain a percentage index of participants' accuracy in recalling these self-assignments, we divided their estimates of the number of times that the self appeared in the clean and dirty groups by 40 and 32 respectively and then multiplied the result by 100. Higher percentages on these indices indicated greater recall accuracy. Overall, participants showed reasonable accuracy ($M = 72.34$, $SD = 38.86$). Interestingly, a paired samples t test revealed that participants were more accurate in recalling the self in the dirty group ($M = 88.83$, $SD = 55.71$) than they were at recalling the self in the clean group ($M = 55.84$, $SD = 35.49$), $t(121) = -7.03$, $p < .001$, $\eta_p^2 = .290$. This greater

accuracy for the self in negative groups is consistent with literature showing that negative information is attended, processed, and recalled more thoroughly than positive information (for a review, see Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001).

The IO Effect and the Moderating Effect of Group Valence

Our predictions regarding the IO effect related to those instances in which participants made errors in the memory recall task and misassigned a target person to the incorrect group. Hence, we focused on this error data in our analyses (for the same approach, see Capozza et al., 2009; Castano et al., 2002; Castelli et al., 2000; Yzerbyt et al., 1995). The mean error rate for the experimental trials was 31.31% ($SD = 11.53$), and there were no outliers on this variable ($\pm 3 SDs$). The mean error rate for the self-assignment trials was 25.51%, $SD = 20.10$. Hence, participants were more accurate in recalling the location of their own code number compared to the location of other people's code numbers. This evidence also demonstrates that participants categorized themselves as members of the correct group just under three-quarters of the time. This is an important check because it implies that participants knew which group they belonged to on the vast majority of the experimental trials.

Note that the inclusion of self-misassignments in our analysis of the IO effect would have biased the data towards an artifactual IO effect because, although it is possible for participants to misassign the self from the in-group to the out-group (i.e., an IO effect), the reverse is not true. In other words, it is not possible for participants to misassign the self from the out-group to the in-group because, by definition, the self is never presented in the out-group in any of the research trials. Hence, we excluded the self-misassignment data from our analyses.

We coded the data from the experimental trials to indicate whether participants had misassigned targets to their in-group (i.e., the group that contained the self) or the out-group (i.e., the group that did not contain the self) as well as whether they misassigned targets to the positive group (i.e., the clean bucket group) or the negative group (i.e., the dirty dustbin group). We then conducted a 2 (group type: in-group/out-group) $\times$ 2 (group valence: positive/negative) repeated measures ANOVA on this data.

Consistent with previous research, there was a significant main effect of group type that indicated an IO effect, $F(1, 121) = 5.23, p = .024, \eta_p^2 = .041$. On average, participants misassigned 10.44 in-group members to the out-group ($SD = 4.44$) but only 9.60 out-group members to the in-group ($SD = 3.98$).

There was also a significant main effect of group valence, $F(1, 121) = 7.05, p = .009, \eta_p^2 = .055$. Participants misassigned significantly more targets from the positive group to the negative group ($M = 10.56, SD = 4.49$) than from the negative group to the positive group ($M = 9.48, SD = 4.12$).

Both of the above effects were qualified by a significant two-way interaction between group type and group valence, $F(1, 121) = 12.10, p = .001, \eta_p^2 = .091$. To investigate this interaction, we tested the effect of group type at each level of group valence. When the in-group was positive and the out-group was negative, participants misassigned significantly fewer targets to the in-group ($M = 4.25, SD = 2.40$) than to the out-group ($M = 5.23, SD = 2.53$), $t(121) = -3.99, p < .001$. However, when the in-group was negative and the out-group was positive, there was no significant difference in misassignments to the ingroup ($M = 5.34, SD = 2.49$) and the out-group ($M = 5.21, SD = 2.73$), $t(121) = .54, p = .590$. Hence, consistent with predictions, the IO

effect was moderated by group valence; it only occurred when the in-group was positive and out-group was negative and not vice versa.² Figure 2 illustrates the two-way interaction between group type and group valence.

Figure 2. Misassignments to the In-Group and Out-Group as a Function of Group Valence

The Predictive Effect of Self-Esteem

After reverse-scoring negatively-worded items, the items in the Rosenberg Self-Esteem Scale showed good internal reliability ($\alpha = .78$), and so we computed an average score for this scale. We also computed an index that represented the IO effect by subtracting the number of misassignments to the in-group from the number of misassignments to the out-group. Larger, positive scores on this index indicated a larger IO effect. Given the moderating effect of group valence, we also computed IO indices separately for misassignments to positive and negative in-groups. We then regressed these IO indices onto self-esteem.

Self-esteem negatively predicted the overall IO effect, $\beta = -.20$, $p = .032$, and the IO effect when the in-group was relatively positive, $\beta = -.21$, $p = .018$. Self-esteem did not predict the nonsignificant IO effect when the in-group was relatively negative, $\beta = -.08$, $p = .377$.³ Hence, consistent with

predictions, the lower people's self-esteem, the more likely they were to misassign targets to the out-group rather than the in-group. Note that this relation was not due to a general lack of accuracy on the part of people with low self-esteem. There was no significant relation between participants' self-esteem and their overall error rate ($p = .454$), and the significant relations between self-esteem and the IO indices remained significant when controlling for error rate in a partial correlation analysis ($ps \leq .039$).

Reaction Time Data

Previous research has found that participants are slower in making misassignments to the in-group than to the out-group (Castano et al., 2002; Yzyerbt et al., 1995). Notably, however, this effect has not always been reliable (Castano, 2004, p. 380), with some studies finding a null effect (Capozza et al., 2009, Study 2) or even a reverse effect (Capozza et al., 2009, Study 1). In the present study, there was no significant difference between these times, $t(121) = 1.37$, $p = .172$.

Perceived Awareness of the Research Hypothesis

After reverse-scoring two negatively-word items, the Perceived Awareness of the Research Hypothesis (PARH) scale had good internal reliability ($\alpha = .77$). Consequently, we computed the average of the four items in order to obtain an index of the extent to which participants believed that they were aware of the research hypotheses during the research. We included this PARH index in correlation analyses with the previously computed indices of the IO effect. The PARH index did not show any significant relations with the overall IO effect ($r = .11$, $p = .237$, $N = 122$), the IO effect in relation to positive in-groups ($r = .09$, $p = .313$, $N = 122$), or the (nonsignificant) IO effect in relation to

negative in-groups ($r = .07, p = .435, N = 122$). Hence, the extent to which participants believed that they were aware of the research hypotheses did not influence the extent to which they showed IO effects. These null effects suggest that the observed IO effects did not result from demand characteristics in our research paradigm.

Discussion

A Better Understanding of the IO Effect

The present research makes a number of important contributions to our understanding of the IO effect. First, the research identified group valence as a moderator of this effect: Participants overexcluded people from the in-group when it was relatively positive but not when it was relatively negative. This moderating effect supports Leyens and Yzerbyt's (1992) motivational explanation of the IO effect, and it provides indirect support for Yzerbyt et al.'s (2000) prediction that the IO effect will be larger for high status groups than for low status groups.

Second, we found that people with low self-esteem showed the strongest IO effect, both in general and when the in-group had a positive valence. These significant associations provide additional support for the motivational explanation: People with low self-esteem have a strong motivation to exclude illegitimate targets from their in-group in order to protect their social identity. Consistent with Brown (1988), these findings add to the literature on social identity theory's self-esteem hypothesis (Abrams & Hogg, 1988; Rubin & Hewstone, 1998) by demonstrating that measures of global personal self-esteem can be used to predict relatively subtle and indirect in-group bias effects such as the IO effect.

Taken together, these results provide relatively direct and conclusive support for Leyens and Yzerbyt's (1992) motivational

explanation of the IO effect. Previous research that has demonstrated the moderating effect of in-group identification is consistent with this motivational explanation (Castano et al., 2002; Knowles & Peng, 2005, Study 3). However, it is also consistent with alternative explanations of the IO effect that are based on the availability of material resources, in-group distinctiveness per se, and differential intergroup familiarity (Corneille et al., 2007; Livingstone et al., 2011; Sherif, 1967). The present research ruled out these alternative explanations by using minimal groups that had no connection with material resources, a fixed and equal size, and no difference in familiarity. In addition, the present research tested predictions regarding group valence and self-esteem that are specific to the motivational explanation and cannot be derived from these alternative explanations.

Finally, the present research makes a significant contribution in this area by providing the first demonstration of the IO effect using minimal groups. Hence, our research shows that the IO effect is a remarkably general phenomenon that extends from long-lived, real-world groups such as linguistic and regional groups to trivial and transient lab-based groups.

Limitations and Directions for Future Research

One disadvantage of our novel recall paradigm is that participants had the opportunity to miscategorize their own code number and falsely believe that they belonged to one group when, in fact, they belonged to the other group. Evidence from the self-assignment trials suggests that this self-misassignment may have occurred in around 25% of the experimental trials. This self-misassignment is likely to have weakened the IO effect because it would result in participants excluding targets from groups that they believed represented their

in-group but that were in fact out-groups. It is a testament to the strength and reliability of the IO effects that we observed that they persisted in spite of this countervailing error-based IO effect.

It is possible that motivational factors may have made this error-based IO effect more prominent when the in-group was negative and the out-group was positive compared to the opposite group valences. This asymmetrical error-based IO effect might then account for the null IO effect that we observed when the in-group was negative. Note that although this error-based IO explanation is also motivational in nature, it is slightly different from the explanation that we put forward. We assumed that the IO effect is nullified for negative in-groups due to a lack of motivation to protect the in-group's status. In contrast, an error-based IO effect explanation assumes that the IO effect is weaker for negative in-groups because people are motivated to miscategorise themselves as members of positive out-groups rather than negative in-groups. Our data from participants' retrospective recall of their group memberships does not support this error-based IO explanation. In particular, participants were more accurate, not less accurate, in recalling the location of their code number in the negative group than in the positive group. Nonetheless, future research should consider ways of controlling for self-misassignment in recall paradigms such as ours.

Given that each group in our paradigm lasted for only 5 seconds, it was not possible to measure participants' identification with their group without disrupting the recall task. Hence, we were unable to ascertain the degree to which our participants identified with their in-groups. However, the fact that we obtained a significant IO effect indicates that some degree of in-group identification must have

occurred. After all, if participants did not classify the people in the same group as them as in-group members at some level, then there would be no basis for the IO effect to occur, and we would not have obtained any significant effect of group type (in-group/out-group). Nonetheless, it remains important for researchers to include measures of identification in future tests of the IO effect in order to arrive at more articulated conclusions on this issue.

We predicted that the IO effect would be weakened, nullified, or even reversed when the in-group had a negative valence and the out-group had a positive valence because, in this situation, group members would not be motivated to keep positive out-group members out of their negative in-group. Our data showed that the IO effect was nullified under these conditions rather than weakened or reversed. However, it is possible that weaker manipulations of group valence may only weaken the IO effect and, conversely, stronger manipulations of group valence may reverse the IO effect. Again, future research should investigate these predictions.

An intriguing question that is raised by our memory recall paradigm is whether IO effects occur due to biases during the attention stage, encoding stage, and/or retrieval stage of information processing. It is informative to consider this issue in some depth because it has implications for the extent to which the present findings may generalize to the more traditional classification paradigm that has been used to investigate the IO effect.

IO effects may occur at a pre-encoding attentional stage of information processing. Participants may attend more to in-group members than to out-group members. This attentional bias may cause a familiarity bias in which people become more familiar with in-group exemplars than

with out-group exemplars (Linville, 1998), and this familiarity bias may then explain the IO effect (Corneille et al., 2007). We do not believe that this attention-based explanation can account for the IO effect that we observed in the present study for three reasons. First, we instructed participants to spend an equal amount of time looking at the code numbers in both groups. Hence, differences in the amount of attention that was paid to in-group and out-group members were minimized and, consequently, unlikely to explain our findings. Second, it was not possible for participants to become more familiar with in-group members than with out-group members in our research because target individuals' changed groups from trial to trial. For example, code number 2 would be an in-group member in some trials but an out-group member in other trials and, as discussed previously, there was no stereotypical association between code numbers and group memberships. Third, the familiarity explanation is not a motivational explanation, and it does not imply the involvement of self-esteem or group valence. Hence, it is unable to explain the results in the current study. Nonetheless, future research should control and/or measure participants' attention towards in-groups and out-groups in order to provide direct evidence that in-group/out-group attentional biases are not responsible for the IO effect (e.g., measuring participants' attention via eye-tracking technology).

IO effects may also occur during the encoding stage of information processing. An encoding-based account is most compatible with the type of IO effects that have been reported in previous classification studies. In this case, people are more likely to incorrectly encode individuals as out-group members rather than as in-group members. Given that several previous studies have demonstrated the occurrence of

this encoding-based IO effect, and that encoding processes must have operated in the present recall task, it is likely that a similar effect occurred in the present research.

Finally, IO effects may also occur during the retrieval stage of a recall task. In this case, even if people correctly encode individuals as in-group members and out-group members, they may nonetheless recall these classifications incorrectly. In other words, they may recall more in-group members as being out-group members than vice versa. This retrieval-based IO effect implies that participants will be slower at misrecalling in-group members than out-group members because they are motivated to be more concerned about making accurate in-group classifications (Castano et al., 2002; Yzyerbt et al., 1995). Contrary to this prediction, our reaction time data revealed no significant difference in the time that it took participants to make in-group and out-group misassignments. Consequently, our results do not support the operation of a retrieval-based IO effect.

The above reasoning and evidence suggests that the IO effect in the present research occurred primarily during the encoding stage. Consequently, our findings are likely to generalize to IO effects that have been demonstrated in the classification paradigm because these effects also occur during the encoding stage. However, future research in this area should measure self-esteem and manipulate group valence in the standard classification paradigm in order to provide a more definitive test of the generalizability of our findings.

Finally, it is also informative to consider how our results relate to other prominent theories of minimal group biases. One such theory is self-anchoring theory (Cadinu & Rothbart, 1996). According to this theory, people generalize the positive views that they hold about themselves (i.e.,

their positive self-esteem) to their in-groups but not to out-groups. Hence, self-anchoring theory predicts a positive relation between people's self-esteem and the degree to which they favour their in-group over out-groups. This prediction stands in contrast to the negative relation that we observed between self-esteem and the IO effect. However, this apparent contradiction can be reconciled after taking into account the different stages at which in-group overexclusion and self-anchoring operate. Chronologically, the IO effect precedes the self-anchoring effect because people must first classify others as in-group members and out-group members before they can project their own positive traits onto in-group members. Hence, although self-esteem may be negatively related to the overexclusion of in-group members, it may nonetheless be positively related to the biased evaluation of these in-group members following classification. Considering these IO and self-anchoring processes together, we might predict that people with high self-esteem would be relatively inclusive but biased regarding their in-group, whereas people with low self-esteem would be relatively exclusive but less biased regarding their in-group.

References

- Abrams, D., & Hogg, M. A. (1988). Comments on the motivational status of self-esteem in social identity and intergroup discrimination. *European Journal of Social Psychology, 18*, 317-334. doi: 10.1002/ejsp.2420180403
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology, 5*, 323-370. doi: 10.1037/1089-2680.5.4.323
- Boccatto, G., Capozza, D., & Falvo, R. (2003). Need for distinctiveness and ingroup overexclusion effect. *Ricerche di Psicologia, 26*, 65-82.
- Brown, J., Collins, R. L., & Schmidt, G. W. (1988). Self-esteem and direct vs. indirect forms of self-enhancement. *Journal of Personality and Social Psychology, 55*, 445-453. doi: 10.1037/0022-3514.55.3.445
- Cadinu, M. R., & Rothbart, M. (1996). Self anchoring and differentiation processes in the minimal group setting. *Journal of Personality and Social Psychology, 70*, 661-677. doi: 10.1037/0022-3514.70.4.661
- Capozza, D., Boccatto, G., Andrighetto, L., & Falvo, R. (2009). Categorization of ambiguous human/ape faces: Protection of ingroup but not outgroup humanity. *Group Processes and Intergroup Relations, 12*, 777-787. doi: 10.1177/1368430209344868
- Castano, E. (2004). In case of death, cling to the ingroup. *European Journal of Social Psychology, 34*, 375-384. doi: 10.1002/ejsp.211
- Castano, E., Yzerbyt, V. Y., Bourguignon, D., & Seron, E. (2002). Who may enter? The impact of ingroup identification on ingroup-outgroup categorization. *Journal of Experimental Social Psychology, 38*, 315-322. doi: 10.1006/jesp.2001.1512
- Castelli, L., Gorrasi, F., & Arcuri, L. (2000). The phenomenon of overexclusion by one's own group during a memory task: A motivational explanation. *Giornale Italiano di Psicologia, 27*, 261-279.
- Corneille, O., Hugenberg, K., & Potter, T. (2007). Applying the attractor field model to social cognition: Perceptual discrimination is facilitated, but memory is impaired for faces displaying evaluatively congruent expressions. *Journal of Personality and Social Psychology, 93*, 335-352. doi: 10.1037/0022-3514.93.3.335
- Dazzi, C., Voci, A., Brambilla, F., & Capozza, D. (1996). Ingroup overexclusion: Stability of the effect. In R. Romaioli, M., Dondi, & C., Dazzi, (Eds). *Experimental and theoretical studies in psychology: Ph.D. students' contributions from the psychology departments of Padua University* (pp. 175-183). Padova, Italy: Cooperativa Libreria Editrice Università di Padova; Italy.
- Devos, T., & Banaji, M. R. (2005). American = White? *Journal of Personality and Social Psychology, 88*, 447-466. doi: 10.1037/0022-3514.88.3.447
- Knowles, E. D., & Peng, K. (2005). White selves: Conceptualizing and measuring a dominant-group identity. *Journal of Personality & Social Psychology, 89*, 223-241. doi: 10.1037/0022-3514.89.2.223
- Leyens, J. -P., & Yzerbyt, V. Y. (1992). The ingroup overexclusion effect: Impact of valence and confirmation on stereotypical information search. *European Journal of Social*

- Psychology*, 22, 549-569. doi: 10.1002/ejsp.2420220604
- Linville, P. W. (1998). The heterogeneity of homogeneity. In J. M. Darley & J. Cooper (Eds.), *Attribution and social interaction: The legacy of Edward E. Jones* (pp. 423-487). Washington, DC: American Psychological Association. doi: 10.1037/10286-008
- Livingstone, A. G., Spears, R., Manstead, A. S. R., & Bruder, M. (2011). The more, the merrier? Numerical strength versus subgroup distinctiveness in minority groups. *Journal of Experimental Social Psychology*, 47, 786-793. doi: 10.1016/j.jesp.2011.03.012
- Pauker, K. B., Weisbuch, M., Ambady, N., Sommers, S. R., Adams, R. B. Jr., & Iyengar, Z. (2009). Not so black and white: Memory for ambiguous group members. *Journal of Personality and Social Psychology*, 96, 795-810. doi: 10.1037/a0013265
- Rosenberg, M. (1965). *Society and the adolescent self-image*. Princeton, NJ: Princeton University Press.
- Rubin, M., & Hewstone, M. (1998). Social identity theory's self-esteem hypothesis: A review and some suggestions for clarification. *Personality and Social Psychology Review*, 2, 40-62. doi: 10.1207/s15327957pspr0201_3
- Rubin, M., Paolini, S., & Crisp, R. J. (2010). A processing fluency explanation of bias against migrants. *Journal of Experimental Social Psychology*, 46, 21-28. doi: 10.1016/j.jesp.2009.09.006
- Sherif, M. (1967). *Group conflict and co-operation*. London: Routledge.
- Tajfel, H., & Turner, J. C. (1979). An integrative theory of intergroup conflict. In S. Worchel & W. G. Austin (Eds.), *The social psychology of intergroup relations* (pp. 33-47). Monterey, CA: Brooks-Cole.
- Yzerbyt, Y., Castano, E., Leyens, J. P., & Paladino, M. (2000). The primacy of the ingroup: The interplay of entitativity and identification. *European Review of Social Psychology*, 11, 257-295. doi: 10.1080/14792772043000059
- Yzerbyt, V. Y., Leyens, J. -P., & Bellour, F. (1995). The ingroup overexclusion effect: Identity concerns in decisions about group membership. *European Journal of Social Psychology*, 25, 1-16. doi: 10.1002/ejsp.2420250102

Footnotes

1. In a part of the procedure that was unrelated to our investigation of the IO effect, participants who were assigned the number 3 were also assigned to a Red Group, and participants who were assigned the number 14 were also assigned to a Blue Group. Participants then responded to nine items that measured their sense of identification with their colour-based group. The identification items were standard, face-valid items that are commonly used in this type of research (e.g., “I feel a strong bond with my group,” “I identify with other people in my group”). Note that the items related only to the colour-based groups and so could not be used to assess identification with the in-groups and out-groups that are discussed in the main text. The target individuals in the memory recall task were also identified as members of the colour-based groups. As in Figure 1, code numbers 1 to 10 were presented in red ink and assigned to the Red Group, and code numbers 11 to 20 were presented in blue ink and assigned to the Blue Group. This part of the procedure was intended to investigate in-group bias, and we found a significant in-group bias effect among high identifiers. However, this effect is not novel, and it is not the focus of the present article. Consequently, we do not discuss it in the main text. It is important to note that colour group membership and code number were counterbalanced across the group type and group valence factors, and they did not interact with the IO effects that are described in the main text ($ps \geq .238$). In addition, group membership and code number had no significant effect on participants’ self-esteem, $t(120) = .86, p = .390$. It is also important to note that it was impossible to investigate the IO effect in relation to colour group membership because participants were unable to misassign participants from one colour group to the other during the

research. Hence, this part of the procedure was entirely separate from our investigation of IO effects. We describe it here solely for the purposes of explaining the use of the colour coding in our procedure.

2. We repeated our ANOVA including participants’ gender as a factor. This factor did not interact with either group type or group valence ($ps \geq .093$).

3. Using Mahalanobis distance, we identified one multivariate outlier on the self-esteem and IO indices ($ps > .001$). However, the exclusion of this outlier did not affect the pattern of significant results that are reported in the main text.