The impact of group membership on punishment versus partner choice

Justin W. Martin, Liane Young, Katherine McAuliffe

Department of Psychology and Neuroscience, Boston College, Chestnut Hill, MA

Word count: 13,534

Correspondence: justin.martin.3@bc.edu (J.W. Martin)

Address: 275 Beacon St, Chestnut Hill, MA 02467

Author note: We are grateful to Masoud Jasbi, Sydney Levine, MH Tessler and Kara Weisman for feedback on earlier versions of this manuscript. This work was supported by the Boston College Virtue Project, the John Templeton Foundation and the Canadian Institute for Advanced Research (Azrieli Global Scholar award). Data and analysis code for this work can be found at <https://osf.io/cjs8q/>. Materials are included in the Supplementary Materials.

Abstract

Human punishment is biased by ingroup favoritism: we tend to punish outgroup members more harshly than ingroup members. Such preferential punishment of outgroup members could reflect an influence of group membership specifically on punishment, or may reflect a general tendency to respond to outgroup members' behavior more harshly, regardless of the type of response used. To investigate this question, we contrasted punishment with the decision to reject a partner, often termed partner choice. In two studies, participants interacted with other players in an incentivized economic game. We assigned participants to groups using a "minimal" groups paradigm (Study 1) or a consequential political position (Study 2). Across both studies, when participants could respond to their partner's behavior with punishment, they punished outgroup members more harshly than ingroup members, replicating past work. We also extend prior work by showing that this difference principally reflects outgroup derogation rather than ingroup love, through the inclusion of neutral individuals in Study 2. In contrast, when participants could respond by either continuing to interact with their current partner or instead reject them and be paired with a new player, participants' decisions were almost completely unaffected by group membership. Thus, group membership has a strong influence on how we punish others, but almost no influence on how we make partner choice decisions. These results shed light on the breadth of influence group membership can have, especially on how we respond to transgressions, and provide insight into the unique psychological processes supporting punishment and partner choice decisions.

Keywords: Punishment, partner choice, intergroup bias, rejection, intentions

Humans are exceptionally sensitive to group membership (Cikara, 2018; Cikara et al., 2011, 2014; Dunham, 2018; McAuliffe & Dunham, 2016; Molenberghs et al., 2014; Riek et al., 2006). Whether someone is a member of our group—in terms of nationality, political party, social club or a host of other categories—influences a variety of behaviors, including with whom we cooperate, empathize and trust, among many others (Cikara, 2018; Dunham, 2018; McAuliffe & Dunham, 2016; Riek et al., 2006; Tajfel & Turner, 1979).

Of recent interest has been the tendency for group membership to influence how we respond when others transgress (e.g. when they harm us or others, break a norm or otherwise behave negatively). Much of this work has focused on a particular way we can respond: by punishing. Punishment is typically thought of as a reaction to a transgression which imposes a cost on the transgressor (Boyd & Richerson, 1992; Clutton-Brock & Parker, 1995; Raihani & Bshary, 2015, 2019). Past research has found a consistent tendency to more harshly punish outgroup members than ingroup members (Baumgartner et al., 2012; Bernhard, Fehr, et al., 2006; Bernhard, Fischbacher, et al., 2006; Delton & Krasnow, 2017; Goette et al., 2012; Gollwitzer & Keller, 2010; Guo et al., 2020; Halevy et al., 2008; Jordan et al., 2014; Kubota et al., 2013; Mendoza et al., 2014; Schiller et al., 2014; Shinada et al., 2004; Valenzuela & Srivastava, 2012; Yudkin et al., 2016) though we also note that work on the black sheep effect has shown the opposite in some cases (Van Prooijen, 2006). However, punishment is not the only way we can respond to others' transgressions. Another option we have is to decide to reject transgressors or stop interacting with them, as in

the case of partner choice (Barclay, 2013, 2016). An open question is whether group membership influences punishment specifically or more broadly influences how we respond to transgressions. Here, we explore this question by contrasting the use of punishment and partner choice in response to unfair intentions from and unfair outcomes produced by ingroup and outgroup members.

Partner choice provides a useful contrast to punishment because it is both similar to and distinct from punishment in important ways. Partner choice is similar to punishment in that it is a method for responding to transgressors, is often a dyadic response (e.g. imposed by one person on another) and involves imposing something negative on the transgressor (a cost in the case of punishment and the removal of a social connection in the case of partner choice). Both responses frequently occur in social contexts, meaning it is plausible that both are influenced by various social factors, including group membership. At the same time, partner choice is importantly different from punishment from both a theoretical and empirical perspective. From a theoretical perspective, punishment has often been considered as a tool with the functional purpose of changing others' behavior (Clutton-Brock & Parker, 1995; Martin & Cushman, 2016; Trivers, 1971), although it has also been studied as a means for costly signaling (Barclay, 2006; Jordan et al., 2016), norm enforcement (Boyd & Richerson, 1992; Fehr & Gächter, 2002; Henrich et al., 2010; Yamagishi, 1986) or fitness leveling (Raihani & Bshary, 2019). In contrast, the theoretical perspective on partner choice focuses on the idea that individuals exist in a biological market (Noe & Hammerstein, 1994), which leads market dynamics to influence a host of factors related to attracting,

choosing and maintaining partners as well as rejecting partners with whom we no longer wish to interact (Barclay, 2013, 2016). Although rejecting a partner following a transgression may impose a cost, just as punishment imposes a cost, a key distinction between partner choice and punishment is that the imposition of a cost is not an inherent or key feature of rejecting a partner. Whereas punishment necessarily involves a cost to the transgressor, rejecting someone may or may not involve a cost for that individual, and any such cost is incidental to the rejection of the transgressor. That is, the principal focus of rejecting a transgressor is on avoiding them, regardless of whether or not it costs that transgressor as a result. Functionally, then, in the context of transgressions partner choice and punishment serve distinct goals—with punishment serving to change others' behavior and partner choice serving as a mechanism for rejecting the current partner and finding a new one (Barclay, 2013, 2016). Given this, we might expect punishment and partner choice to be used differently in response to bad behavior.

Indeed, initial empirical evidence suggests that punishment and partner choice are used in different ways (Liddell & Kruschke, 2014; Martin & Cushman, 2015). In these studies, participants were given the chance to respond to a partner's intention to be fair or selfish and to the outcome that was produced in an economic game. When participants had the opportunity to punish (i.e. take money away from the partner or not), participants were influenced relatively equally by the transgressor's intent and the outcome that they produced. In contrast, when participants had the opportunity to make a partner choice decision (i.e. to reject the partner and not play with them on the next

round of the game), participants were influenced much more by the transgressor's intent and much less by the outcome produced. Thus, punishment and partner choice are differently sensitive to intentions and outcome, with partner choice more sensitive to others' intentions and punishment more sensitive to outcomes caused. This distinction between punishment and partner choice both theoretically and empirically makes partner choice ideal for testing the breadth of the influence group membership has on how people respond to others' transgressions.

Past work provides hints about how group membership might influence partner choice decisions. One possibility is that partner choice, like punishment, is sensitive to group membership. From a functional perspective, partner choice is often thought about as serving a group regulatory function (Baumard et al., 2013; Bednarik et al., 2014; van der Lee et al., 2017), which would imply a sensitivity to group membership. Indeed, the mere availability of partner choice as an option leads group members to behave more cooperatively (Bednarik et al., 2014; Wu et al., 2016), and group competition leads to increased rejection of those who contribute less to the group (Sääksvuori et al., 2011). Most relevantly, feedback about the aggregate amount of cooperation from ingroup and outgroup members leads to higher rates of cooperation in a prisoner's dilemma, but only when participants can choose their partner compared to when their partner is randomly assigned (Aksoy, 2015). Critically, these studies provide only indirect evidence about the influence of group membership on how we respond to others' transgressions, in that they either do not manipulate group membership (Bednarik et al., 2014; Sääksvuori et al., 2011; Wu et al., 2016) or do not provide

participants with information about the target's past behavior (Aksoy, 2015).

Furthermore, the most relevant of these studies investigates how we *choose* a partner based upon group membership (Aksoy, 2015) rather than how we decide whether or not to *reject* them based upon prior negative behavior. Selecting partners amongst alternatives and deciding whether to reject a current partner in favor of alternatives are distinct aspects of partner choice psychology (Barclay, 2013, 2016). Thus, it is plausible that group membership may influence one aspect of partner choice (i.e. choosing partners) but not other aspects (i.e. rejecting partners).

Beyond these results, other indirect evidence comes from work focusing on how group membership influences those who experience partner choice, which finds a strong sensitivity to group boundaries. Being excluded hurts more when it comes from ingroup members (Sacco et al., 2014), especially when the group is based on permanent features (e.g. gender) as opposed to impermanent features (e.g. t-shirt color) (Wirth & Williams, 2009). Being excluded and then-reincluded by a group leads to greater risk-taking on behalf of that group (Dahl et al., 2019), and children who are excluded by their ingroup are subsequently more anxious and more adherent to ingroup norms (Watson-Jones et al., 2016). Among those who most strongly identify with their group, rejection by that group leads to greater endorsement of pro-group actions (Gómez et al., 2011). These results demonstrate that group membership influences how people respond to being rejected, providing some reason to believe that group membership would also play a role in the decision to reject or not. In this sense, it is

plausible that group membership will influence both punishment and partner choice decisions.

An alternative possibility is that group membership may not influence both partner choice and punishment, and will only have an influence on punishment. There are a few reasons to consider this possibility. First, as discussed above, the most relevant prior work has found sensitivity to group membership when *choosing* a partner (Aksoy, 2015) and has not focused on the influence of group membership when *rejecting* a partner after a transgression. Given that these are distinct aspects of partner choice psychology (Barclay, 2013, 2016), it is reasonable to consider whether group membership might exert less of an influence when rejecting a partner than when choosing a partner.

Second, at a fundamental level, partner choice in the context of another's transgression is about deciding to either end that interaction or continue it. This requires a pre-existing relationship, not necessarily in the rich sense of a close, personal relationship, but at least in the thinner sense of having interacted with this person in a joint endeavor. Presumably, those with whom we have this type of relationship are more likely to share membership with us in at least some groups, compared to those with whom we do not have this type of relationship. Thus, group membership may be a less relevant feature when deciding whether to end or continue a relationship—because shared group membership is a consistent feature of those we are considering ending a relationship with, group membership is not a pivotal factor in this decision. Analogously, when deciding what kind of computer to buy, I do not consider

whether the computer has a built-in wireless adapter; this feature is present in most (though not all) computers and so it does not even enter the decision-making process. Potentially then, group membership is held constant enough in those with whom we have interacted previously that this feature is less relevant when deciding whether to end or continue that interaction. Critically, such an argument applies less well to punishment – we can just as easily punish those with whom we have and have not interacted. Indeed, punishing outgroup members may be important, especially in the context of intergroup conflict. Consider political disagreements—transgressions by one's own party member may be tolerated, whereas transgressions by other party members are often strongly condemned and cause for retribution.

Third, when in a partner choice mindset, group membership may become uncertain and therefore exert less of an impact. That is, while contemplating whether to continue interacting with someone or not, social boundaries are clearly fluid—I am no longer committed to continuing our relationship and am considering ending it. In some sense, this is an opportunity to start fresh, and so past entanglements—including group membership—may also become more fluid. This fluidity may mean that group membership is uncertain enough that it exerts a reduced influence on whether or not we decide to reject this individual.

Thus, there are good reasons to consider the possibility that group membership may exert less of an influence on whether we choose to reject a partner for their past behavior than on how we punish them.

Here, we seek to explore these two possibilities, focusing on the degree to which punishment and partner choice show equivalent or distinct sensitivity to group membership. Answering this question provides insight into how these two distinct responses are made. More broadly, these results will shed light on how pervasive of an influence group membership has on how we respond to others' transgressions—if decisions to punish and decisions to end an interaction are both sensitive to group membership, that would provide initial evidence that group membership has a broad impact in this context (although there may yet be ways of responding to transgressions that are not influenced by group membership). However, if we find that the influence of group boundaries is specific to punishment (or partner choice), that would suggest a more limited role for group membership in how we respond when others transgress.

In answering this question, we employ a paradigm developed in prior work to differentiate punishment and partner choice (Martin & Cushman, 2015). This game asks participants to engage in either punishment or partner choice in response to an allocator's intent to be fair or selfish and the kind of outcome produced—a fair or selfish division of the stake. Specifically, participants played an online version of a stochastic economic game in which they were asynchronously paired with other participants who pre-committed to behavior in the game. In this game, a Decider allocated \$1 between themselves and the Responder by choosing between two options, one of which was more likely to produce a fair outcome (an even split of the money between the players) and the other of which was more likely to produce an unfair outcome (all \$1 going to the Decider). Critically, either option could produce either outcome. Thus, the Decider's

intent—to be fair versus selfish—is dissociable from whether there is a fair or selfish outcome. Participants played as the Responder and were told the Decider's choice of option and which outcome was produced. In the Punishment condition, the Responder could take bonus money away from the Decider's payoff at a cost to themselves. In the Partner Choice condition, the Responder decided whether, for another round of the game, they would like to switch to a new Decider or stick with the same Decider. Prior work (Martin & Cushman, 2015) found that responses in the Partner Choice condition were heavily influenced by the Decider's intent to be fair or selfish and very little by the outcome produced. In contrast, responses in the Punishment condition were relatively equally influenced by the Decider's intent and the outcome produced. Thus, this paradigm effectively dissociates punishment and partner choice decisions, making it a perfect testbed for an investigation of the role of group membership.

We move beyond this prior work by having participants play this economic game as part of distinct teams. In line with some past work (Goette et al., 2012), we use two complementary ways of determining team membership. In Study 1, we employ a variant of the “minimal” groups paradigm (Dunham, 2018; Tajfel, 1970). Participants are assigned to a team based upon which of two words they first identify in a simple word search (“owl” or “cat”) (Figure 1). This paradigm has the virtue of putting participants into groups about which they know almost nothing, allowing us to investigate the influence of the abstract structure of thinking about groups on punishment and partner choice decisions. Of course, a lack of sensitivity to group boundaries in such a case may not be particularly informative—such a group carries no

real-world significance and so a lack of sensitivity may reflect the relatively weak nature of this manipulation. To address this concern, in Study 2 we assign participants to groups based upon their position on a consequential political position—their stance on abortion. Amongst a series of other demographic items, participants reported whether they are pro-life, pro-choice or hold no position (data from these participants was not included). Participants were then assigned to a group based upon this position. This procedure has the advantage of invoking highly consequential boundaries which are frequently relevant in daily life. Thus, this group assignment has the potential to intensify the influence of group membership on punishment and partner choice decisions. On the other hand, this procedure makes use of pre-existing differences across participants, and so is a less clean test of the influence of group boundaries compared to Study 1. Using both approaches allows us to extrapolate beyond the weakness of each in order to paint a clearer picture of the influence of group boundaries on how we respond to others' transgressions. Regardless of the method of assigning participants to groups, choosing to reject a partner changed only the individual that the participant would play with for the next round, not their group membership. This was done so as to match the Partner Choice and Punishment conditions as much as possible (given that participants could not change their partner's group membership in the Punishment condition).

Beyond providing a replication of Study 1 and extending our results to real-world groups, Study 2 included additional methodological improvements. In Study 1, participants interacted with only ingroup and outgroup members, whereas in Study 2 we

also included a neutral individual, allowing us to examine whether ingroup bias is reflective of ingroup love or outgroup derogation (relative to a baseline individual). While Study 1 employed punishment responses on a scale from -\$0.30 to \$0.30, Study 2 used a binary measure (-\$0.30 or \$0), allowing us to analyze data from the Punishment and Partner Choice conditions within the same regression model. Finally, in order to follow up on some analyses in Study 1, Study 2 included a measure of cognitive reflection. Specifically, we included a revised 7-item version of the Cognitive Reflection Test (CRT) (Pennycook & Rand, 2020). This test is designed to assess the degree to which participants make decisions by reflecting/deliberating or using their gut reaction/first instinct.

Study 1

Data availability

Data and analysis scripts for all studies can be found at <https://osf.io/cjs8q/>. For both studies, we report all measures, manipulations and exclusions.

The pre-registration for this study can be found at

<http://aspredicted.org/blind.php?x=j92vh7>. All materials for Study 1, including instructions to participants, can be found in the Supplementary Material.

Method

Participants

Participants (n = 802) were recruited via Amazon Mechanical Turk and completed an online survey in exchange for a small payment ($\leq$ \$2.00) with an

opportunity to earn a bonus payment. Sample size was determined and pre-registered before any data analysis. Participants made decisions in an asynchronous economic game involving other participants, and their data was excluded from analysis on the basis of pre-registered criteria (see Procedures below for description of the economic game, comprehension questions and attention check questions): Comprehension questions at the beginning of the task, answers to attention check questions at the end of the task, reaction time during the main task (mean reaction time less than three standard deviations below the overall mean, after log transformation) and reporting not being a native speaker of English. These criteria yielded a final sample of 659 participants (17.8% excluded; final sample 48% male, 51% female, <0.5% other, <0.5% preferred not to answer). The percent excluded was very similar across our response type conditions (Punishment: 11.3%; Partner Choice: 14.9%). All participants were United States residents.

A sensitivity analysis of our main analysis, the comparison between the Punishment and Partner Choice conditions in proportion of participants sensitive to group membership, indicated that our design is able to detect a proportion difference in sensitivity to group membership as small as 7.7% with 80% power at 5% error rate (critical $z = 1.96$). This analysis was performed using G*power version 3.1 (Faul et al., 2009).

Procedures

Participants interacted with other participants in an asynchronous stochastic economic game modeled on past work (Cushman et al., 2009; Martin & Cushman,

2015). Before playing the game, participants were assigned to one of two minimal groups and subsequently played two rounds of the game, one with another player on their own team and one with another player on the other team. No deception was used in either of our studies—participants were truthfully matched with other participants, and the decisions participants made had a real impact on the other player's bonus payment (and their own bonus payment). All procedures were approved by the Boston College Institutional Review Board.

Group assignment. After providing consent, participants were told they would perform a task that would determine which team they were on before engaging in an economic game with other players. On the next screen, participants were presented with a 3x3 word search (see Figure 1) and told to enter the first three-letter word they saw into a free response box. They were told that the experiment would assign them to a team based on their response. The experiment would proceed only if the participant answered either “cat” or “owl”, ignoring case. Participants' responses were distributed between the two teams (403 Cat, 256 Owl). Participants were then directed to a screen informing them that there are two teams, TEAM CAT or TEAM OWL, and that “people on the same team tend to share many personality and cognitive traits”. This was done so as to promote group identification. Based on the word entered, participants were then informed they had been assigned to either TEAM CAT or TEAM OWL and selected one of five potential cat or owl avatars to represent them throughout the task.

Figure 1. The design of Studies 1 and 2. **A:** In both the Punishment and Partner Choice conditions, participants interacted with others players in an economic game for real stakes. Participants played as both Responders and Deciders. In the Punishment condition, Responders could add to or subtract from the Deciders payoff. In the Partner Choice condition, Responders decided whether, for another round of the game, they would like to play with the same Decider or a new Decider (from the same group). **B:** The group assignment mechanism in both studies. Study 1 employed a minimal-groups paradigm. Participants were instructed to enter the first 3-letter word they found in the word search depicted (answers were “OWL” or “CAT”). Participants were assigned to a group based upon their choice and selected an avatar to represent them in the game. Study 2 employed real-world groups. Amongst other demographic items collected at the beginning of the study, participants reported their stance on the issue of abortion. Participants were assigned to a group based upon their response. **C:** Following their responses in the economic game, participants rated their identification with the groups used in the study. In both studies, participants identified more with their ingroup than their outgroup. In Study 2, the neutral group was identified with less than the ingroup but more than the outgroup.

Stochastic economic game. Participants were then given instructions about the economic game they would play. This game involved two roles: The Decider and the Responder. On each trial, the Decider allocated \$1 between themselves the Responder by choosing between two options. Option A had a $\frac{2}{3}$ chance of giving all money to the Decider and a $\frac{1}{3}$ chance of giving \$0.50 to the Decider and \$0.50 to the Responder. Option B had a $\frac{2}{3}$ chance of giving \$0.50 to the Decider and \$0.50 to the Responder and a $\frac{1}{3}$ chance of giving all money to the Decider. Thus, Option A is more likely to yield an unfair division of the money than Option B, though both options can yield either outcome. In this way, the Decider's intent to be fair or selfish is easily inferable from the option they choose, though they can produce a fair or selfish outcome regardless of their choice.

Responders responded to the Decider's choice of option and outcome of their choice, with the response available varying by condition (between-participants). In the Partner Choice condition, participants were told that they would play another round of the game and decided whether they wanted to stay with the same Decider for the second round or to instead switch to a new Decider. If the participant decided to switch, this new Decider would be from the same group as their current partner; switching thus did not allow participants to play with a new group member. Choosing to stay or switch cost the Responder nothing. In the Punishment condition, participants had the opportunity to add or subtract up to \$0.30 from the Decider's payoff, at no cost to themselves. We chose for punishment to be costless to enact so as to mirror the cost-free nature of the decision in the Partner Choice condition. In both response type

conditions, we employed the strategy method: participants were asked to commit to a response for all four possible combinations of the Decider's choice and the outcome produce (Option A, fair split; Option A, unfair split; Option B, fair split; Option B, unfair split; note that "fair" and "unfair" were never used in the game itself). This method achieved two things: (1) we gained four times the amount of data than using "hot" decisions, and (2) we could ask for participants' decisions without informing them of the Decider's decision or outcome on that round, so that the decisions across rounds were not influenced by how prior Deciders played or which outcomes resulted.

Participants played this game twice, once with an ingroup member and once with an outgroup member. Whether participants played with an ingroup or outgroup member first was randomized across participants.

After the instructions but before participating in the game, participants answered seven comprehension questions regarding the main features of the game¹. If any of these questions was answered incorrectly, participants were shown the correct answer and then asked the question again on a subsequent page. If participants answered incorrectly on this 2nd try, they were counted as getting the question incorrect.

Participants were allowed to complete the task, but their data were subsequently excluded if any of these questions were answered incorrectly. This resulted in exclusion of data for 45 participants in the Partner Choice condition (11.3%) and 60 participants in the Punishment condition (14.9%).

¹ One of these questions was about the number of rounds played in each game (e.g. with each other player). In the Punishment condition, the correct answer is 1. However, the survey incorrectly indicated that answers other than 2 were incorrect. Thus, this question was not used in excluding participants' data from the Punishment condition.

Rating group identification. After playing the game with both an ingroup and outgroup player, participants then completed identification ratings for their ingroup and outgroup. We include these ratings after participants' economic game participation to avoid demand effects which may occur when rating identification immediately following group assignment. Specifically, participants rated their agreement with four statements drawn from prior work (Leach et al., 2008) about each group: "I like TEAM [team]", "Members of TEAM [team] are good", "I feel attached to TEAM [team]" and "I identify with TEAM [team]". The scale used went from "Strongly disagree" to "Strongly agree" on 0 to 100-point slide scale. The order in which participants rated their ingroup and outgroup was randomized across participants. Reliability across these four items was high for both Ingroup (Cronbach's $\alpha = 0.92$, 95% CI = 0.90–0.92) and Outgroup (Cronbach's $\alpha = 0.82$, 95% CI = 0.80–0.84) members, so an aggregate rating for each participant was created for both Ingroup and Outgroup members by averaging responses to all four items for that group.

Playing as Decider. Participants were then told they would play this game twice more, but now in the role of Decider. Principally this provided us with Deciders that we could subsequently match to Responders, but this also provided additional potentially useful data. Participants were presented with the same set of instructions as they read earlier, but not presented from the Decider's perspective. They again answered seven comprehension questions, though these were presented only once, and participants were not screened based upon their answers. Finally, participants played as the Decider twice, once with an ingroup player and once with an outgroup player.

Attention check questions, demographics and debrief. Next, participants answered a series of attention check questions regarding attentiveness to the study (e.g. “Overall, how much attention did you pay to this study while you were taking it?”), which were used in data quality screening. Finally, participants answered a series of optional demographic variables and were debriefed about the goals and hypotheses of the study.

Analysis approach. In Study 1, our primary analysis investigated the sensitivity of participants in the Punishment and Partner Choice conditions to the Decider’s intent, the outcome of the die roll, and the Decider’s group membership. Because responses across our two response type conditions are on different scales, the straightforward use of regression models to compare these conditions is not possible. Specifically, responses in the Punishment condition used an interval scale (from -\$0.30 to \$0.30 by \$0.10), whereas responses in Partner Choice condition used a binary scale (stay with Decider, switch to new Decider), preventing us from analyzing these conditions in the same regression model. Thus, as specified in our pre-registered analysis plan and consistent with past work (Martin & Cushman, 2015), we analyze data by classifying each participant as being sensitive to intentions, outcomes, and group membership (non-mutually exclusive) and then comparing these percentages across the Punishment and Partner Choice conditions.

Participants were classified as sensitive to a feature (intentions, outcomes, and group membership) if, holding the other two features constant, their response differed based upon the levels of that feature (ignoring the direction of that difference). For

instance, participants were classified as sensitive to intentions if, holding outcome and group membership constant, their responses differed as a function of the Decider's intent. So, if a participant responded differently when the ingroup Decider chose Option A and produced a fair outcome versus when the ingroup Decider chose Option B and produced a fair outcome (possibilities differing in intent but not outcome or group membership), they would be classified as sensitive to intentions. And if a participant responded differently when the ingroup Decider chose Option B and produced a fair outcome versus when the outgroup Decider chose Option B and produced a fair outcome (possibilities differing in group membership but not intent or outcome), they would be classified as sensitive to group membership, for example. Note that this classification does not consider the direction of difference. In addition, we note that classification of each of these features (intentions, outcomes, and group membership) was conducted independently (and not mutually exclusively), so participants could be classified as sensitive to a subset of these features (e.g. intentions and outcomes but not group membership), all of them, or none of them.

To additionally explore our data, responses within the Punishment and Partner Choice conditions were analyzed using regression models. In these analyses, the main fixed effects included group membership (ingroup versus outgroup), intent (selfish versus fair), and outcome (selfish versus fair), as well as all possible interactions between these variables. Additional fixed effects were included depending upon the goal of the model and are noted where appropriate.

Responses in the Punishment condition were interval data (from -0.30 to 0.30 by 0.10), and so linear mixed-effects regression was used, implemented in R using the lme4 package (D Bates et al., 2014). These models included a random intercept for each participant as well as random slopes for group membership, intent and outcome. For these models, we use a model comparison approach and evaluated the importance of predictors using Likelihood Ratio Tests. We began by comparing the full model to one dropping the highest-level term (e.g. the 3-way interaction in a model with 3 main factors), asking whether this highest-level term improved model fit. If so, we proceeded to inspect this model further. If not, this term was removed and we next compared a model including all terms at the lower level (e.g. all 2-way interactions in this example) to a model without those terms. We continued until one level of interaction terms is significant or until no interactions are significant.

Responses in the Partner Choice condition were binary data (stay versus switch), and so logistic regression was used. Mixed-effects models were not used to analyze data in this condition, as these models indicated that 0 variance was associated with the random intercept for participants. Looking at means across participants (ignoring our three factors), 271 out of 327 (83.0%) participant means were 0.5, with the other eight potential values each accounting for less than 15 participants (<4.3%). Thus, standard logistic regression was used to explore these results.

When analyzing participants' decisions when playing as the Decider, responses are binary (Option A versus Option B), and so we model these data using mixed-effects

logistic regression, including a random intercept for participants and again taking a model comparison approach.

Results and Discussion

Did participants identify more with ingroup members?

We first asked whether our group manipulation was successful: did participants identify more with their ingroup members than outgroup members? To this end, for each participant we calculated their ingroup identification (aggregating across all four questions) and their outgroup identification. We found that participants identified with their ingroup significantly more than their outgroup (Figure 1; Ingroup mean: 62.5, SEM: 0.91; Outgroup mean: 41.0, SEM: 0.71; paired t -test(658) = 24.37, $p < 0.001$, Cohen's $d = 1.34$, 95% Confidence Interval of mean difference = 19.78–23.25). This was true both in the Punishment condition (Ingroup mean: 63.7, SEM: 1.33; Outgroup mean: 41.5, SEM: 1.06; paired t -test(331) = 17.59, $p < 0.001$, Cohen's $d = 1.36$, 95% CI of mean difference = 19.80–24.79) and the Partner Choice condition (Ingroup mean: 61.2, SEM: 1.24; Outgroup mean: 40.4, SEM: 0.95; paired t -test(326) = 16.86, $p < 0.001$, Cohen's $d = 1.32$, 95% CI of mean difference = 18.30–23.14), with no difference across conditions in identification for either ingroup members (Welch two-sample t -test(654.63) = 1.41, $p = 0.16$, Cohen's $d = 0.11$, 95% CI of mean difference = -1.01–6.12) or outgroup members (Welch two-sample t -test(649.55) = 0.69, $p = 0.49$, Cohen's $d = 0.05$, 95% CI of mean difference = -1.82–3.78). Thus, our group manipulation successfully led participants to identify with their own group over the alternative group and did not have a different effect across our two response type conditions.

Did group membership have a different influence on punishment and partner choice?

We next asked whether group membership had a differential influence on punishment and partner choice decisions. We found that it did: those in the Partner Choice condition were less sensitive to group membership (Punishment: 29.8%; Partner Choice: 12.2%; $\chi^2(1) = 29.57$, $p < 0.001$, 95% CI of proportion difference = 11.2%–24.0%), more sensitive to intentions (Punishment: 40.4%; Partner Choice: 52.3%; $\chi^2(1) = 8.96$, $p = 0.003$, 95% CI of proportion difference = 4.1%–19.8%) and less sensitive to outcomes (Punishment: 66.0%; Partner Choice: 55.4%; $\chi^2(1) = 7.34$, $p = 0.007$, 95% CI of proportion difference = 3.0%–18.3%) than those in the Punishment condition. Thus, we found a significant difference in how sensitive participants are to group membership in the Punishment and Partner Choice conditions.

This was also true when accounting for differences across conditions in participants who were “insensitive” to the main features of our task (Figure 2): i.e., those who were not sensitive to intentions, outcomes or group membership. That is, consistent with past work (Martin & Cushman, 2015), we found a significant difference across conditions in these participants (Punishment: 26.5%; Partner Choice: 4.0%; $\chi^2(1) = 62.72$, $p < 0.001$). Because these participants lower the proportion of participants exhibiting sensitivity to any of our factors differently across conditions, we compared classification of participants across conditions both with and without these participants included, as specified in our pre-registered analysis plan. When removing these

insensitive participants, we found overall similar results. Those in the Partner Choice condition continued to be less sensitive to group membership (Punishment: 39.3%; Partner Choice: 12.7%; $\chi^2(1) = 51.3$, $p < 0.001$, 95% CI of proportion difference = 19.1%–34.1%) and less sensitive to outcomes (Punishment: 88.5%; Partner Choice: 57.6%; $\chi^2(1) = 62.3$, $p < 0.001$, 95% CI of proportion difference = 23.7%–38.0%), though these analyses revealed no difference in sensitivity to intentions across conditions (Punishment: 54.1%; Partner Choice: 54.4%; $\chi^2(1) < 0.001$, $p = 1$, 95% CI of proportion difference = -9.1%–8.3%). Thus, regardless of how insensitive participants were treated, we found support for our key result: greater sensitivity to group membership (and less sensitivity to outcomes) amongst those in the Punishment condition. In contrast, differences in sensitivity to intentions changed depending upon the presence or absence of insensitive participants.

These analyses revealed a difference in how sensitive punishment and partner choice were to group membership: those in the Punishment condition were more sensitive to whether the person they are interacting with was an ingroup or outgroup member compared to those in the Partner Choice condition. However, these analyses did not reveal the direction of that difference, i.e. whether those in the Punishment condition punished ingroup members more or less. We next turned to mean responding in each condition to explore this question.

Figure 2. *Participants' sensitivity to intentions, outcomes and group membership across the Punishment and Partner Choice conditions in Study 1.* Plotted is the percentage of participants in the Punishment and Partner Choice conditions whose responses exhibited sensitivity to intentions, outcomes, and group membership (coded in a non-mutually exclusive manner). Also depicted is the percentage of participants who were not sensitive to any features of the task (i.e. who responded uniformly).

Punishment condition. Confirming the above results, we found that responses in the Punishment condition were sensitive to intentions, outcomes, and group membership (Figure 3). Specifically, while we did not find a 3-way interaction between intentions, outcomes, and group membership ($LRT\ X^2(1) = 0.01, p = 0.93$), we did find that 2-way interactions improved model fit ($LRT\ X^2(3) = 12.53, p = 0.006$). Inspecting this model further, we found a significant interaction between intentions and outcomes ($B = -0.09, SE = 0.03, t = -3.48, p < 0.001, 95\% CI = -0.14 - -0.04$), but no interaction between group membership and either intentions ($B = -0.02, SE = 0.03, t = -0.65, p = 0.52, 95\% CI = -0.07 - 0.03$) or outcomes ($B = -0.004, SE = 0.03, t = -0.14, p = 0.89, 95\% CI = -0.05 - 0.04$). We also found main effects of intentions ($B = -0.41, SE = 0.05, t = -8.58, p < 0.001, 95\% CI = -0.50 - -0.31$), outcomes ($B = -0.73, SE = 0.05, t = -15.98, p < 0.001, 95\% CI = -0.81 - -0.63$), and, critically, group membership ($B = 0.06, SE = 0.02, t = 3.81, p < 0.001, 95\% CI = 0.03 - 0.09$). Thus, consistent with the classification results above, those in the Punishment condition were more sensitive to outcomes than intentions, judging by standardized beta values, and were additionally sensitive to group membership.

Looking at overall mean responding (Figure 3), we found that the interaction between intentions and outcomes was explained by a somewhat greater influence of intentions when the outcome was unfair (Fair intent, unfair outcome: $-0.05, SEM = 0.01$; Selfish intent, unfair outcome: $-0.14, SEM = 0.01$) than when the outcome was fair (Fair intent, fair outcome: $0.10, SEM = 0.01$; Selfish intent, fair outcome: $0.02, SEM = 0.01$). In addition, we found an overall influence of group membership that did not interact with

intentions or outcomes: participants were more lenient when interacting with an ingroup member (mean: -0.01, SEM = 0.01) than an outgroup member (mean: -0.03, SEM = 0.01).

We additionally explored the boundary conditions for and moderators of this influence of group membership on punishment (see Supplemental Material). We found that the overall pattern of results, including the main effect of group membership on punishment, was unchanged if insensitive participants were removed. We also found that those who most strongly identified with their ingroup were most lenient toward ingroup members, while continuing to find an overall influence of group membership on punishment. We found that participants were more lenient toward ingroup members when the first partner they encountered was an ingroup member than when this partner was an outgroup member. However, we again found an overall influence of group membership on punishment even when this influence of order was accounted for. Finally, lenience toward ingroup members was greatest amongst those who were most attentive (as measured by first-try performance on comprehension questions). Critically, we continued to find an overall influence of group membership on punishment even when this influence of attentiveness was accounted for. Thus, while we found a number of factors that increase leniency in punishment of ingroup members, none of these factors eliminated the overall effect of group membership on punishment.

Partner Choice condition. In the case of partner choice, we found a very different pattern. In particular, we found no main effect of group membership (OR = 1.09, SE = 0.15, $z = 0.56$, $p = 0.57$, 95% CI of OR = 0.82–1.45) and no interactions

involving group membership (all OR < 1.20 and > 0.85, all $z < |0.56|$, all $p > 0.52$). We did find a significant interaction between intentions and outcomes (OR = 2.05, SE = 0.29, $z = 2.47$, $p = 0.01$, 95% CI of OR = 1.18–3.69) as well as relatively equal main effects of intentions (OR = 0.06, SE = 0.15, $z = -19.89$, $p < 0.001$, 95% CI of OR = 0.04–0.07) and outcomes (OR = 0.05, SE = 0.15, $z = -20.9$, $p < 0.001$, 95% CI of OR = 0.04–0.06). When insensitive participants were removed, we no longer found an interaction between the influence of intentions and outcomes, though we continued to find relatively equal main effects of both factors and no interactions with or main effect of group membership (see Supplemental Materials). Looking at overall mean responding (Figure 3), we found that the interaction between intentions and outcomes was driven by a greater tendency to stay with the current Decider based upon the presence of an unfair outcome when they had a fair intent (Fair intent, fair outcome: $P(\text{stay with current Decider}) = 0.97$, SEM = 0.01; Fair intent, unfair outcome: 0.50, SEM = 0.02) than when they had an unfair intent (Unfair intent, fair outcome: 0.54, SE of proportion = 0.02; Unfair intent, unfair outcome: 0.07, SEM = 0.01).

As with punishment, we additionally explored influence of a number of factors on the relationship between group membership and punishment (see Supplemental Material). When removing insensitive participants, we found that group membership still had no significant impact on partner choice. However, we no longer found an interaction between the influence of intentions and outcomes. We also found that those who most strongly identified with their ingroup took their partner's intention into account to a greater extent. However, we did not find that strength of identification led group

membership to influence partner choice decisions. We found a complicated pattern based upon the order in which participants encountered ingroup and outgroup Deciders. Participants weighted outcomes more heavily for the first Decider they encountered, staying if this decider caused a fair outcome and switching if they caused an unfair outcome. Importantly, we continued to find no overall influence of group membership, nor any interaction between group membership and intentions or outcomes. Finally, similar to the relationship between ingroup identification and the influence of intentions, we found that those who were most attentive took intentions into account to a greater extent. Critically, this factor did not influence the contribution of group membership, and we continued to find no overall influence of group membership. Thus, none of these factors led group membership to have an overall influence on partner choice decisions.

Decisions as the Decider. Finally, because all participants played the game as a Decider after their decisions as a Responder were collected, we investigated how participants' decisions as a Decider were influenced by group membership and whether they might be subject to punishment or partner choice (see Supplemental Material). To summarize, we found that participants were more likely to choose the fair option when the Responder was an ingroup member than an outgroup member, in both response conditions, but also that participants were more likely to choose the fair option when the Responder had the option to punish the participant versus when the Responder had the option to switch away from the participant. And, the tendency to be fairer toward ingroup Responders was accentuated when individual differences in ingroup identification were taken into account.

Figure 3. Mean responses in the Punishment and Partner Choice conditions in Study 1. Plotted are participants' mean responses across the Intention, Outcome, and Group membership conditions for the Punishment (Panel A) and Partner Choice (Panel B) conditions. Error bars are SEM. Note that possible punishment values ranged from - \$0.30 to \$0.30; we plot a restricted range here for visualization purposes.

Study 2

Data availability

The pre-registration for this study can be found at <http://aspredicted.org/blind.php?x=56xa9a>. All materials for Study 2, including instructions to participants, can be found in the Supplementary Material.

Method

Participants

Participants (n = 865) were recruited via Amazon Mechanical Turk and completed an online survey in exchange for a small payment ($\leq \$2.00$) with an opportunity to earn a bonus payment. Sample size was determined and pre-registered before any data analysis. The task used was a modified version of the asynchronous economic game used in Study 1 (see below for description). Exclusionary criteria included all criteria from Study 1 (performance on comprehension questions; answers to attention check questions at the end of the task; reporting not being a native speaker of English; mean reaction time less than three standard deviations below the overall mean, after log transformation), with the additional criteria of reporting no position on the political issue used to assign participants to a group. These criteria yielded a final sample of 600 participants (23.2% excluded; final sample 47% male, 52% female, $<0.5\%$ other, $<0.5\%$ preferred not to answer). The percent excluded was very similar across our response type conditions (Punishment: 22.7%; Partner Choice: 23.8%). All participants were United States residents.

A sensitivity analysis of our main analysis, the model comparing responses from participants in the Punishment and Partner Choice conditions, indicated that our design has 80% power with 5% error rate to detect an effect (Odd's Ratio) as small as 0.62 or 1.60 for a 2-way interaction (our main result) and as small as 0.11 or 8.17 for a 4-way interaction. This analysis was performed using the *simr* package (Green & Macleod, 2016) in R, with 1000 bootstrap samples.

Procedures

Participants interacted with other participants in a modified version of the economic game used in Study 1. Before playing the game, participants were assigned to one of three groups based on their position on a political issue. As in Study 1, no deception was used. All procedures were approved by the Boston College Institutional Review Board.

Group assignment. After providing consent, participants answered a series of standard demographic questions, plus a question about their stance on a political issue. Specifically, they were asked “Regarding the issue of abortion, how much do you identify as Pro-Choice vs. Pro-Life?”, with “Pro-Choice”, “Pro-Life” and “No position” as possible answers. They also rated how strongly they supported this position on a scale from 1 = “Not strongly at all” and 7 = “Very strongly”. Participants and the other players they interacted with were subsequently identified with an image corresponding to their choice (Figure 1).

Stochastic economic game. Participants played the same stochastic economic game from Study 1, with a few modifications. First, participants now played this game

three times, once with a player who expressed a pro-life position, once with a player who expressed a pro-choice position and once with a player who expressed no position. The order in which participants interacted with these players was counterbalanced across participants. Second, only two response options were available in the Punishment condition. Specifically, we analyzed participants' punishment decisions in Study 1 to select the two options which we most frequently chosen. We found that the most frequently chosen option was to not punish (41% of responses) and that the next most frequently chosen option was maximum punishment (removing \$0.30; 27% of responses). Thus, we selected these two amounts as the punishment options available in Study 2.

Participants again answered a series seven comprehension questions, using the same procedure as in Study 1.

Rating group identification. After playing the game with all three players, participants then rated their identification with each type of player, using the same scale as in Study 1. The order in which player types were rated was randomized across participants. As in Study 1, these ratings had high reliability (Ingroup: 0.91, 95% CI = 0.89–0.91; Outgroup: 0.92, 95% CI = 0.91–0.93; Neutral: 0.89, 95% CI = 0.88–0.90) and so an aggregate measure was calculated for each participant for each group by averaging the 4 individual ratings for that group.

Playing as Decider. Participants then played the game three more times, now playing as the Decider. Participants were provided with the same set of instructions but now described from the Decider's perspective. Participants again answered a series of

7 comprehension questions, presented only once and not used in screening data quality. Participants played the game once with each type of player, with order counterbalanced across participants.

CRT, attention check questions and debrief. Participants next completed the 7-item CRT (Pennycook & Rand, 2020). Then they completed six questions designed to assess the degree to which the other player's unfair choice and the unfair outcome were perceived as negative. For each type of player, participants were asked "When the partner who identified as [type] chose the unfair option (Option A), how negative did this make you feel?" and "When the partner who identified as [type] received the entire \$1, how negative did this make you feel?", with both questions asked on 7-point scales anchored at 1 = "Not at all negative" and 7 = "Extremely negative". Finally, participants completed the same set of attention check questions as in Study 1 and were debriefed.

Analysis approach. In Study 2, our primary analyses used mixed-effects logistic regression, as specified in our pre-registration. Responses to engage in either punishment or partner choice were coded as 1, with not engaging in punishment or partner choice coded as 0. The main fixed effects included response type (punishment versus partner choice), group membership (3-level factor: ingroup, outgroup versus neutral), intent (selfish versus fair), and outcome (selfish versus fair), as well as all possible interactions between these variables. Additional fixed effects were included depending upon the goal of the model and are noted where appropriate. We include a random intercept for participants. Including random slopes generated significant convergence issues. Inspection of these models and the distribution of our response

variable indicated that some cells of our design had very few responses. For instance, when paired with an Ingroup Decider, in the case where the Decider had a fair intent and produced a fair outcome, only 8/300 responses (0.3%) were to reject the Decider, and 13/300 responses (0.4%) were to engage in punishment. Similarly, when paired with a Neutral Decider, in the same case, only 9/300 responses (0.3%) were to reject the Decider and 22/300 responses (0.7%) were to engage in punishment. Typical procedures to address convergence issues were not able to overcome this issue of rare events. Because mixed-effects models only including random intercepts can yield inflated type 1 error rates (Barr, Levy, Scheepers, & Tily, 2013; but see Douglas Bates, Kliegl, Vasishth, & Baayen, 2015; Matuschek, Kliegl, Vasishth, Baayen, & Bates, 2017), we additionally model our results using standard logistic regression (without random effects) while performing Firth bias correction (Firth, 1993), which reduces small-sample bias and yields definite estimates. In all cases, these models either found the same pattern of statistical results or yielded a larger number of significant fixed effects, suggesting that our random effects models were not inflating type 1 error rates. Models with random effects were implemented in R using the lme4 package (D Bates et al., 2014), while models without random effects implemented in R using the base package, with Firth bias correction implemented using the brglm2 package (Kosmidis, 2020). As in Study 1, we use a model comparison approach and evaluate our data using Likelihood Ratio Tests.

Additionally, in order to facilitate comparison to the results of Study 1, we classified each participant as being sensitive to intentions, outcomes and group

membership (non-mutually exclusive) and then compared these percentages across the Punishment and Partner Choice conditions (see Supplemental Material). We used the same approach as in Study 1, classifying participants as sensitive to a feature (intentions, outcomes and group membership) if, holding the other two features constant, their response differed based upon the levels of that feature (ignoring the direction of that difference).

When analyzing participants' decisions when playing as the Decider, responses are binary (Option A versus Option B) and so we model these data using mixed-effects logistic regression, including a random intercept for participants and again taking a model comparison approach.

Results and Discussion

Did participants identify more with their ingroup?

We again began by asking whether our group manipulation was successful. To this end, for each participant we calculated their ingroup identification (aggregating across all four questions), their outgroup identification and their identification with neutral individuals. We then compared these ratings within both the Punishment and Partner Choice conditions. We found that participants identified with their ingroup significantly more than their outgroup or neutral individuals (Figure 1; see Supplemental Material for full results). This was true in both the Punishment and Partner Choice conditions. We found no difference across conditions in preference for these three groups. Thus, our group manipulation successfully led participants to identify more with

their own group over the alternative group as well as neutral individuals and our manipulation did not have a different effect across our two response type conditions.

Did group membership have a different influence on punishment and partner choice?

Turning to our main results (Figure 4), we again found a stronger influence of group membership on punishment. Running a regression including predictors for response type, group membership, intentions, and outcomes, we did not find that the 4-way interaction between these variables improved model fit (LRT $\chi^2(2) = 0.38$, $p = 0.83$) but did find a significant contribution of 3-way interactions (LRT $\chi^2(7) = 14.37$, $p = 0.04$). Inspecting this model further (Model 1; see Supplemental Table 1 for full model results), we found a significant interaction between group membership, intentions, and outcomes (Ingroup vs. Outgroup: OR = 2.66, SE = 0.39, $z = 2.52$, $p = 0.01$, 95% CI = 1.25–5.78; Ingroup vs. Neutral: OR = 1.45, SE = 0.42, $z = 0.89$, $p = 0.37$, 95% CI = 0.64–3.30; Neutral vs. Outgroup: OR = 1.84, SE = 0.37, $z = 1.65$, $p = 0.10$, 95% CI = 0.90–3.81), as well as significant 2-way interactions between group membership and response type (Ingroup vs. Outgroup: OR = 2.13, SE = 0.18, $z = 4.30$, $p < 0.001$, 95% CI = 1.51–3.00; Ingroup vs. Neutral: OR = 1.25, SE = 0.18, $z = 1.22$, $p = 0.22$, 95% CI = 0.87–1.79; Neutral vs. Outgroup: OR = 1.70, SE = 0.17, $z = 3.12$, $p = 0.002$, 95% CI = 1.22–2.38), and group membership and outcomes (Ingroup vs. Outgroup: OR = 0.52, SE = 0.19, $z = -3.39$, $p < 0.001$, 95% CI = 0.35–0.75; Ingroup vs. Neutral: OR = 0.92, SE = 0.21, $z = -0.39$, $p = 0.69$, 95% CI = 0.61–1.38; Neutral vs. Outgroup: OR = 0.56,

SE = 0.19, $z = -3.14$, $p = 0.002$, 95% CI = 0.39–0.80). We also found a marginal interaction between response type, intentions, and outcomes (OR = 1.83, SE = 0.32, $z = 1.89$, $p = 0.06$, 95% CI = 0.98–3.43), and 2-way interactions between response type and intentions (OR = 0.14, SE = 0.28, $z = -6.98$, $p < 0.001$, 95% CI = 0.08–0.25), and outcome and intentions (OR = 0.43, SE = 0.31, $z = -2.76$, $p = 0.006$, 95% CI = 0.23–0.77). We found main effects for (Ingroup vs. Outgroup: OR = 2.18, SE = 0.10, $z = 8.10$, $p < 0.001$, 95% CI = 1.81–2.64; Ingroup vs. Neutral: OR = 1.29, SE = 0.10, $z = 2.54$, $p = 0.01$, 95% CI = 1.06–1.58; Neutral vs. Outgroup: OR = 1.68, SE = 0.09, $z = 5.76$, $p < 0.001$, 95% CI = 1.41–2.01), response type (OR = 0.34, SE = 0.19, $z = -5.76$, $p < 0.001$, 95% CI = 0.23–0.49), intentions (OR = 8.11, SE = 0.16, $z = 13.45$, $p < 0.001$, 95% CI = 6.03–11.12) and outcomes (OR = 27.49, SE = 0.16, $z = 20.94$, $p < 0.001$, 95% CI = 20.4–37.9). Thus, we found significant relationships between response type and group membership, as well as interactions between these variables and the influence of intentions and outcomes. These results were unchanged if logistic regression with Firth bias correction was employed (Model 8, Supplemental Table 2). We also found a similar pattern of results using the classification approach employed in Study 1 (see Supplemental Material), though we note that the difference in sensitivity to group membership between the Punishment and Partner Choice conditions was marginal when insensitive participants were included ($p = 0.06$). In total, these results demonstrate that group membership had a strong influence on punishment but not on partner choice.

To unpack these interactions, we inspected mean responding. First, we focused on the interaction between response type and group membership. This was driven by greater sensitivity to group membership amongst those in the Punishment condition (Ingroup = 0.29, SEM = 0.01; Neutral = 0.33, SEM = 0.01; Outgroup = 0.42, SEM = 0.01) than those in the Partner Choice condition (Ingroup = 0.44, SEM = 0.01; Neutral = 0.45, SEM = 0.01; Outgroup = 0.47, SEM = 0.01). Indeed, analyzing only data from the Partner Choice condition, we found no 3-way interaction (LRT $\chi^2(2) = 4.04$, $p = 0.13$) and while did we find that 2-way interactions improved model fit (LRT $\chi^2(5) = 14.45$, $p = 0.01$), we found no interactions with group membership (all OR < 1.11 and > 0.72, $p > 0.20$). We also found no difference in judgment of ingroup and neutral partners (OR = 1.21, SE = 0.11, $z = 0.69$, $p = 0.49$, 95% CI = 0.87–1.34) or neutral and outgroup partners (OR = 1.17, SE = 0.11, $z = 1.44$, $p = 0.15$, 95% CI = 0.95–1.45), though we did find a weak difference in judgment of ingroup and outgroup partners (OR = 1.26, SE = 0.11, $z = 2.14$, $p = 0.03$, 95% CI = 1.02–1.57). Thus, these results replicated the findings from Study 1. We again showed that ingroup members were favored over outgroup members (i.e. were punished less) to a much greater extent when making punishment decisions. Study 2 also extended these findings by including a neutral individual, allowing these results to speak to whether responses in our paradigm were driven more by preferences for ingroup members or dislike of outgroup members. We found that this interaction between response type and group membership was driven by greater punishment of outgroup members relative to either ingroup members or neutral individuals. Thus, our group differences appear to reflect outgroup derogation rather

than ingroup love. Finally, we found almost no sensitivity to group membership amongst those in the Partner Choice condition, except a weak overall difference between judgment of ingroup versus outgroup partners. These results largely replicated the findings of Study 1, in which group membership had no influence on those in the Partner Choice condition.

Next, looking at the interaction between group membership, intentions, and outcomes, we found that the difference in responding between the outgroup Decider and the two other players was largest when the outcome was fair (*Fair intent, fair outcome*; Ingroup = 0.04, SEM = 0.01; Neutral = 0.05, SEM = 0.01; Outgroup = 0.12, SEM = 0.01; *Unfair intent, fair outcome*; Ingroup = 0.27, SEM = 0.02; Neutral = 0.29, SEM = 0.02; Outgroup = 0.37, SEM = 0.02) than when the outcome was unfair (*Fair intent, unfair outcome*; Ingroup = 0.44, SEM = 0.02; Neutral = 0.47, SEM = 0.02; Outgroup = 0.49, SEM = 0.02; *Unfair intent, unfair outcome*; Ingroup = 0.71, SEM = 0.02; Neutral = 0.75, SEM = 0.02; Outgroup = 0.79, SEM = 0.02). Thus, the influence of good outcomes was discounted when produced by outgroup members.

Finally, we also found a marginal interaction between response type, intentions and outcomes. Here, we see that outcomes exerted a much stronger influence (and intentions exerted a much weaker influence) on punishment decisions (*Fair intent, fair outcome* = 0.10, SEM = 0.01; *Unfair intent, fair outcome* = 0.18, SEM = 0.01; *Fair intent, unfair outcome* = 0.48, SEM = 0.02; *Unfair intent, unfair outcome* = 0.63, SEM = 0.02) than on partner choice decisions (*Fair intent, fair outcome* = 0.04, SEM = 0.01; *Unfair intent, fair outcome*: 0.44, SEM = 0.02; *Fair intent, unfair outcome* = 0.45, SEM = 0.02;

Unfair intent, unfair outcome: 0.87, $SEM = 0.01$). This pattern replicated the results of Study 1, in which intentions exerted a greater influence on partner choice decisions.

As in Study 1, we explored the influence of a number of additional variables on this overall pattern of results, especially the interaction between group membership and response type (see Supplemental Material). First, we found that these results were unchanged if insensitive participants were excluded, with the exception that we found additional interactions between group membership, intentions and outcomes, and between response type, intentions and outcomes. Similarly, we continued to find an interaction between response type and group membership when including ingroup identification ratings in our models. However, we now additionally found 2-way interactions between ingroup identification and group membership, outcomes and intentions, as well as a marginal interaction between ingroup identification and response type. Including a predictor for which Decider participants interacted with first yielded no influence of Decider order on the influence of group membership. We continued to find an interaction between response type and group membership, though we found additional interactions between Decider order and intentions and Decider order and outcomes. Including a predictor for attentiveness (as measured by first-try performance on comprehension questions), we found an interaction between attentiveness, outcomes and group membership; attentiveness, response type and outcomes; and between attentiveness, response type and intentions. However, attentiveness did not influence the interaction between response type and group membership, which remained significant even when the role of attentiveness was accounted for. Finally,

including CRT performance in our models yielded a marginal interaction between CRT performance, response type and group membership. Inspecting means, this interaction reflected a greater tendency to differentiate outgroup partners from ingroup and neutral partners amongst those higher on CRT performance, relative to those lower on CRT performance. In addition, those higher on CRT performance were less likely to engage in punishment overall and more likely to reject the Decider, relative to those lower on CRT performance. We note, however, that strong conclusions should not be drawn from these results, given their marginal nature. We also found interactions between CRT performance, response type and outcomes and between CRT performance, response type and intentions. However, we continued to find an interaction between response type and group membership when the influence of CRT performance was accounted for. Thus, while these additional factors influenced the role that other factors play in some cases, the interaction between response type and group membership was consistently significant and only ever modulated in the case of the marginal interaction with CRT performance. In total, we consistently found support for the idea that group membership has a strong influence on punishment and has almost no influence on partner choice decisions.

Finally, as in Study 1, we investigated how participants' decisions as a Decider were influenced by group membership and whether they might be subject to punishment or partner choice (see Supplemental Material). To summarize, we found that participants were most likely to choose the fair option when the Responder was an ingroup member, slightly less likely to choose the fair option when paired with a neutral

individual, and least likely to choose the fair option when paired with an outgroup member, in both response conditions. In Study 2 we no longer found that response type influenced these decisions. We again found that the influence of Responder group membership was stronger when individual differences in ingroup identification were taken into account, such that those highest on ingroup bias were more likely to favor ingroup Responders over outgroup or neutral Responders.

Figure 4. Mean responses in the Punishment and Partner Choice conditions in Study 2. Plotted are participants' mean responses across the Intention, Outcome, and Group membership conditions for the Punishment (Panel A) and Partner Choice (Panel B) conditions. Error bars are SEM.

	Study 1	Study 2
Is punishment or partner choice more sensitive to group membership?	Punishment	Punishment
... even when removing insensitive participants?	Yes	Yes
... even when accounting for individual differences in ingroup preference?	Yes	Yes
... even when accounting for order in which partners were interacted with?	Yes	Yes
... even when accounting for individual differences in attentiveness?	Yes	Yes
... even when accounting for individual differences in CRT performance?	--	Yes

Table 1. Summary of results across Study 1 and Study 2.

General discussion

We investigated how group membership influences responses to transgressions by contrasting punishment with the decision to reject this partner, often termed partner choice (Barclay, 2013, 2016). In two studies, participants interacted with other players in an economic game for real stakes. We assigned participants to groups on the basis of both a “minimal” groups paradigm (Study 1) and a consequential political position (Study 2). Across both studies, we find that when participants can respond to the other player’s behavior with punishment, they punish outgroup members more harshly than ingroups members, replicating past work. In contrast, when they can either reject this player or continue interacting with them, participants are almost completely insensitive to group membership. Thus, we find evidence that group membership has a strong influence on how we punish others, but almost no influence on how we make partner choice decisions.

Implications of these results for understanding the psychology underlying partner choice and punishment

Our results show that punishment is much more sensitive to group membership than partner choice. This distinction in how punishment and partner choice decisions are made accords well with past empirical work demonstrating differences in how these two responses are made (Liddell & Kruschke, 2014; Martin & Cushman, 2015). Here, we find differential sensitivity to group membership. We find this pattern both using a minimal groups paradigm (Study 1) and when assigning participants to groups on the

basis of a consequential political position (Study 2). Finding similar results using two distinct ways of determining group membership demonstrates that our results are robust to the type of group assignment mechanism employed, giving greater confidence that the pattern we find may extend to a variety of types of social groups.

Our results are also consistent with theoretical work suggesting different functional purposes for punishment (Barclay, 2006; Boyd & Richerson, 1992; Clutton-Brock & Parker, 1995; Fehr & Gächter, 2002; Henrich et al., 2010; Jordan et al., 2016; Martin & Cushman, 2016; Trivers, 1971; Yamagishi, 1986) and partner choice (Barclay, 2013, 2016; Noe & Hammerstein, 1994). However, the idea that punishment and partner choice are distinct both behaviorally and functionally does not explain the particular pattern we observe here. Why might punishment be sensitive to group membership in a way that partner choice is not? We offer two potential explanations for why partner choice is less sensitive to group membership than punishment. First, rejection of an interaction partner because of a transgression typically occurs in the context of an ingroup relationship. Those we have interacted with previously are likely to be ingroup members on a number of dimensions, more so than those we have not interacted with. Thus, among the factors we might consider when deciding to reject this person or not, their group membership may not feature as prominently since it is a dimension across which there is less variability. In contrast, because we can punish both those with whom we have and have not interacted previously, punishment decisions may be much more sensitive to this person's group membership. Second, when deciding whether or not to reject someone after a transgression, group membership

may already be uncertain enough that its influence is reduced. That is, once someone harms you, you may be uncertain about their group standing – it seems unlikely that an ingroup member would harm you. Their group membership may therefore be uncertain and open to revision. This fluidity regarding group membership may mean that this factor plays less of a role on whether or not we decide to reject this individual. Of course, the present data do not adjudicate between these potential explanations, or any other explanation, only demonstrating a difference between punishment and partner choice in sensitivity to group membership. Indeed, the experimental context used here does not allow us to investigate the influence of group membership relative to a variety of other possible factors. Future work should continue to explore the psychological mechanisms responsible for the difference between punishment and partner choice in sensitivity to group membership demonstrated here.

Relationship with past work on partner choice and punishment

Beyond this difference in sensitivity to group membership, we replicate past work demonstrating a difference between punishment and partner choice in sensitivity to an agent's intentions and the outcomes they cause (Liddell & Kruschke, 2014; Martin & Cushman, 2015). In particular, across both Study 1 and Study 2 we find that intentions exert a stronger influence on partner choice decisions compared to punishment, and outcomes exert a weaker influence on partner choice decisions compared to punishment decisions. We thus extend these prior results from undergraduate samples (Liddell & Kruschke, 2014; Martin & Cushman, 2015) to a more representative and

heterogeneous sample of participants (on Amazon Mechanical Turk). In addition, we note that the sensitivity of punishment and partner choice to intentions and outcomes is much greater than the sensitivity to group membership (see Figures 3 and 4), and that overall effect size for group membership was relatively weak. Thus, while we find reliable differences in sensitivity to group membership between punishment and partner choice, this influence is best understood in the context of much stronger sensitivity to intentions and outcomes.

Beyond replicating these prior results, we also find some differences between our results and those of prior studies. Specifically, we found an unexpectedly weak influence of intentions in the Partner Choice condition compared to past work using this paradigm (Martin & Cushman, 2015). We note that we find a much greater influence of intentions on partner choice decisions than punishment decisions, which is consistent with the relative weighting of intent across these two responses. But at an absolute level, intentions mattered less in our studies than in past work, especially for partner choice decisions. Two features of our data help explain this pattern. First, in both Study 1 and Study 2, participants who were more “attentive”—in the sense of getting more comprehension questions correct on their first try—incorporated intent into their partner choice decisions to a greater extent. Furthermore, in Study 2 we find that those who performed better on the Cognitive Reflection Test—a measure of participants’ tendency to deliberate and reflect—incorporated intentions into their partner choice decisions to a greater extent. Combined, these results suggest that the reduced overall influence of intentions can be explained by differences across samples in attention paid

to the study or participants' willingness to more deeply consider how best to perform in this task. Whereas prior studies have used in-lab undergraduate samples (Liddell & Kruschke, 2014; Martin & Cushman, 2015), participants in our studies were recruited online through Amazon Mechanical Turk. These prior undergraduate samples may have been more invested in the task than our online participants. We note, however, that we did exclude participants who were unable to answer comprehension checks questions correctly, so this difference in attentiveness would have to be an influence that did not impact participants' understanding of the task. Finally, we note that including attentiveness and CRT performance in our analyses did not change the influence of group membership.

Implications of these results for understanding group bias

Our results have two implications of our understanding of group bias more broadly. First, our results allow us to investigate the influence of group membership not just on how participants respond to others' transgressions, but also on their decisions to allocate resources in a fair versus unfair way. This is because participants played our economic game both as Responder and as a Decider (to facilitate participants interacting with real partners). Comparing decisions as Responder and a Decider, we find that sensitivity to group membership as a Responder predicted sensitivity to group membership as a Decider: Those who favored ingroup members as a Responder were more fair toward ingroup members as a Decider. Thus, we find that those who are

sensitive to group membership when responding to others' behavior are also incorporate this information when allocating resources.

Second, Study 2 allowed us to investigate the nature of the influence group membership has on punishment, through the inclusion of a neutral partner. Most work on the influence of group membership on punishment has not included a neutral or unaffiliated individual (Baumgartner et al., 2012; Bernhard, Fehr, et al., 2006; Bernhard, Fischbacher, et al., 2006; Delton & Krasnow, 2017; Goette et al., 2012; Guo et al., 2020; Halevy et al., 2008; Jordan et al., 2014; Kubota et al., 2013; Mendoza et al., 2014; Shinada et al., 2004; Valenzuela & Srivastava, 2012; Yudkin et al., 2016). Only one prior study, to our knowledge, has included such a manipulation (Schiller et al., 2014). This study found both harsher punishment of outgroup members and more lenient punishment of ingroup members relative to neutral individuals, though the effect for outgroup punishment was three times as large (Schiller et al., 2014). Here, we find largely consistent results, in that we observe significantly greater punishment of outgroup members than neutral individuals (and ingroup members). Contrary to prior results, however, we find no difference in punishment between neutral individuals and ingroup members. Of course, our results were directionally the same as this prior work: ingroup members were punished less than neutral individuals, although the difference was not significant. And, in the cells in which the difference in punishment for ingroup and neutral individuals was the smallest (i.e. when a fair outcome was produced), punishment was near floor, potentially artificially reducing this difference. Given this, future work should continue to explore the robustness of differences in punishment

between ingroup members and neutral individuals. What seems clear, in contrast, is that group bias in the context of punishment is driven more by outgroup derogation than ingroup love. This result thus adds to the broader conversation on whether group bias reflects ingroup love versus outgroup hate (Cikara & Van Bavel, 2014; Halevy et al., 2008; Hewstone et al., 2002). We add to other work demonstrating that group bias in the case of punishment may be driven to a greater extent by outgroup hate.

Implications of these results for understanding the impact of group membership on second-party punishment

Beyond their implications for understanding partner choice, our results bear upon the sensitivity of group membership in cases of second-party punishment in particular. Though past work has robustly demonstrated greater punishment of outgroup members (Baumgartner et al., 2012; Bernhard, Fehr, et al., 2006; Bernhard, Fischbacher, et al., 2006; Delton & Krasnow, 2017; Gollwitzer & Keller, 2010; Halevy et al., 2008; Jordan et al., 2014; Kubota et al., 2013; Schiller et al., 2014; Shinada et al., 2004; Valenzuela & Srivastava, 2012; Yudkin et al., 2016), context matters: greater punishment of outgroup members occurs most consistently in cases where the punisher is a third party (Baumgartner et al., 2012; Delton & Krasnow, 2017; Gollwitzer & Keller, 2010; Halevy et al., 2008; Jordan et al., 2014; Schiller et al., 2014; Shinada et al., 2004; Yudkin et al., 2016) and less consistently when the punisher is the victim (a second party; Kubota et al., 2013; McAuliffe & Dunham, 2017; McLeish & Oxoby, 2011; Mendoza et al., 2014; Stagnaro et al., 2018; Valenzuela & Srivastava, 2012). That is, while greater second-

party punishment of outgroup members has been demonstrated (Kubota et al., 2013; Valenzuela & Srivastava, 2012), it is sometimes either weaker than third-party punishment of group members or even non-existent (McAuliffe & Dunham, 2017; McLeish & Oxoby, 2011; Mendoza et al., 2014; Stagnaro et al., 2018). Our results fit into the former category, demonstrating that second-party punishment is indeed sensitive to group membership. Furthermore, our results accord well with a recent proposal on why this pattern exists. Specifically, Stagnaro and colleagues (2018) argue that ingroup preferences exert a greater influence in less strategic contexts, like the Dictator Game, in which the target cannot influence the participant's outcome following their response. In contrast, in cases where the target can respond and influence the participant's outcome, like the Ultimatum Game, group membership exerts a smaller influence. To the extent that Deciders could not impact participants' payoff once punishment had either been administered or withheld, our paradigm is less strategic and so our results support their argument.

Limitations

A few limitations of our results are worth mentioning. First, we have interpreted the difference between punishment and partner choice in sensitivity to group membership as reflecting asymmetries in the psychology underlying these responses. However, a potential alternative interpretation of our results is that, in the context of our task, those in the Partner Choice condition are simply less attentive or discriminating than those in the Punishment condition. This might manifest as less sensitivity to group

membership. We view this interpretation as unlikely because participants in the Partner Choice condition were not less sensitive to all features of our task. Indeed, partner choice decisions were more sensitive to the Decider's intent in both Study 1 and Study 2. These participants were also sensitive to what outcome was produced. Thus, these participants were selectively less sensitive than those in the Punishment condition to group boundaries.

Second, in order to equate decisions in the Punishment and Partner Choice conditions as closely as possible, punishing was costless for Responders. Thus, it is possible that a different pattern of results would emerge were punishment in our economic game to be costly to Responders. We view this possibility as unlikely, given that a host of prior work has consistently shown an influence of group membership on punishment (Baumgartner et al., 2012; Bernhard, Fehr, et al., 2006; Bernhard, Fischbacher, et al., 2006; Delton & Krasnow, 2017; Halevy et al., 2008; Jordan et al., 2014; Kubota et al., 2013; Schiller et al., 2014; Shinada et al., 2004; Stagnaro et al., 2018; Valenzuela & Srivastava, 2012; Yudkin et al., 2016), in many cases when punishment was costly.

Conclusion

Group-based psychology pervades our social life. Here, we have investigated the extent to which, in the context of others' transgressors, group membership influences only how we punish or also influences how we reject partners. Using two different manipulations of group membership, we find that group membership influences

punishment decisions much more than it influences partner choice decisions. These results shed light on the unique psychologies underlying decisions to punish and decisions to reject a partner. More broadly, our results also expand our understanding of the role that groups play in our social decision-making, suggesting that this influence may be lessened depending upon how one chooses to respond to bad behavior.

References

- Aksoy, O. (2015). Effects of Heterogeneity and Homophily on Cooperation. *Social Psychology Quarterly*, 78(4), 324–344. <https://doi.org/10.1177/0190272515612403>
- Barclay, P. (2006). Reputational benefits for altruistic punishment. *Evolution and Human Behavior*, 27(5), 325–344. <https://doi.org/10.1016/j.evolhumbehav.2006.01.003>
- Barclay, P. (2013). Strategies for cooperation in biological markets, especially for humans. *Evolution and Human Behavior*, 34(3), 164–175. <https://doi.org/10.1016/j.evolhumbehav.2013.02.002>
- Barclay, P. (2016). Biological markets and the effects of partner choice on cooperation and friendship. *Current Opinion in Psychology*, 7, 33–38. <https://doi.org/10.1016/j.copsyc.2015.07.012>
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. <https://doi.org/10.1016/j.jml.2012.11.001>
- Bates, D, Maechler, M., Bolker, B., & Walker, S. (2014). *lme4: Linear mixed-effects models using Eigen and S4. R package version 1.1-7*.
- Bates, Douglas, Kliegl, R., Vasishth, S., & Baayen, H. (2015). *Parsimonious Mixed Models*. 1–27. <https://doi.org/arXiv:1506.04967>
- Baumard, N., André, J.-B., & Sperber, D. (2013). A mutualistic approach to morality: the evolution of fairness by partner choice. *The Behavioral and Brain Sciences*, 36(1), 59–78. <https://doi.org/10.1017/S0140525X11002202>
- Baumgartner, T., Götze, L., Gögler, R., & Fehr, E. (2012). The mentalizing network

- orchestrates the impact of parochial altruism on social norm enforcement. *Human Brain Mapping*, 33(6), 1452–1469. <https://doi.org/10.1002/hbm.21298>
- Bednarik, P., Fehr, K., & Semmann, D. (2014). Costs for switching partners reduce network dynamics but not cooperative behaviour. *Proceedings of the Royal Society B: Biological Sciences*, 281(1792), 20141661–20141661. <https://doi.org/10.1098/rspb.2014.1661>
- Bernhard, H., Fehr, E., & Fischbacher, U. (2006). Group Affiliation and Altruistic Norm Enforcement. *The American Economic Review*, 96(2), 217–221. <http://www.jstor.org/stable/30034645>
- Bernhard, H., Fischbacher, U., & Fehr, E. (2006). Parochial altruism in humans. *Nature*, 442(7105), 912. <https://doi.org/10.1038/nature04981>
- Boyd, R., & Richerson, P. (1992). Punishment allows the evolution of cooperation (or anything else) in sizable groups. *Ethology and Sociobiology*, 195, 171–195. <http://www.sciencedirect.com/science/article/pii/016230959290032Y>
- Cikara, M. (2018). Pleasure in response to outgroup pain as a motivator of intergroup aggression. *Atlas of Moral Psychology*, 193–200. <http://nrs.harvard.edu/urn-3:HUL.InstRepos:37093456%0Ahttp://search.ebscohost.com/login.aspx?direct=true&db=psych&AN=2017-57514-013&site=ehost-live>
- Cikara, M., Botvinick, M. M., & Fiske, S. T. (2011). Us Versus Them: Social Identity Shapes Neural Responses to Intergroup Competition and Harm. *Psychological Science*, 22(3), 306–313. <https://doi.org/10.1177/0956797610397667>
- Cikara, M., Bruneau, E., Van Bavel, J. J., & Saxe, R. (2014). Their pain gives us

pleasure: How intergroup dynamics shape empathic failures and counter-empathic responses. *Journal of Experimental Social Psychology*, 55, 110–125.

<https://doi.org/10.1016/j.jesp.2014.06.007>

Cikara, M., & Van Bavel, J. (2014). The neuroscience of intergroup relations: An integrative review. *Perspectives on Psychological Science*.

Clutton-Brock, T., & Parker, G. (1995). Punishment in animal societies. *Nature*, 373, 209–216.

http://www.anth.ucsb.edu/faculty/gurven/anth169/punishment_in_animal_societies.pdf

Cushman, F., Dreber, A., Wang, Y., & Costa, J. (2009). Accidental outcomes guide punishment in a “trembling hand” game. *PloS One*, 4(8), e6699.

<https://doi.org/10.1371/journal.pone.0006699>

Dahl, E., Niedbala, E. M., & Hohman, Z. P. (2019). Loving the Group That Denies You First: Social Identity Effects of Ostracism Before Inclusion. *Personality and Social Psychology Bulletin*, 45(2), 284–299. <https://doi.org/10.1177/0146167218784901>

Delton, A. W., & Krasnow, M. M. (2017). The psychology of deterrence explains why group membership matters for third-party punishment. *Evolution and Human Behavior*, August. <https://doi.org/10.1016/j.evolhumbehav.2017.07.003>

Dunham, Y. (2018). Mere Membership. *Trends in Cognitive Sciences*, 22(9), 780–793. <https://doi.org/10.1016/j.tics.2018.06.004>

Faul, F., Erdfelder, E., Buchner, A., & Lang, A.-G. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior*

- Research Methods*, 41(4), 1149–1160. <https://doi.org/10.3758/BRM.41.4.1149>
- Fehr, E., & Gächter, S. (2002). Altruistic punishment in humans. *Nature*, 415(6868), 137–140. <https://doi.org/10.1038/415137a>
- Firth, D. (1993). Bias Reduction of Maximum Likelihood Estimates. *Biometrika Trust*, 80(1), 27–38.
- Goette, L., Huffman, D., & Meier, S. (2012). The Impact of Social Ties on Group Interactions: Evidence from Minimal Groups and Randomly Assigned Real Groups. *American Economic Journal: Microeconomics*, 4(1), 101–115. <https://doi.org/10.1257/mic.4.1.101>
- Gollwitzer, M., & Keller, L. (2010). What You Did Only Matters if You Are One of Us. *Social Psychology*, 41(1), 20–26. <https://doi.org/10.1027/1864-9335/a000004>
- Gómez, Á., Morales, J. F., Hart, S., Vázquez, A., & Swann, W. B. (2011). Rejected and excluded forevermore, but even more devoted: Irrevocable ostracism intensifies loyalty to the group among identity-fused persons. *Personality and Social Psychology Bulletin*, 37(12), 1574–1586. <https://doi.org/10.1177/0146167211424580>
- Green, P., & Macleod, C. J. (2016). SIMR: An R package for power analysis of generalized linear mixed models by simulation. *Methods in Ecology and Evolution*, 7(4), 493–498. <https://doi.org/10.1111/2041-210X.12504>
- Guo, R., Ding, J., & Wu, Z. (2020). How intergroup relation moderates group bias in Third-Party Punishment. *Acta Psychologica*, 205(July 2019), 103055. <https://doi.org/10.1016/j.actpsy.2020.103055>

Halevy, N., Bornstein, G., & Sagiv, L. (2008). “In-group love” and “out-group hate” as motives for individual participation in intergroup conflict: A new game paradigm: Research article. *Psychological Science*, 19(4), 405–411.

<https://doi.org/10.1111/j.1467-9280.2008.02100.x>

Henrich, J., Ensminger, J., McElreath, R., Barr, A., Barrett, C., Bolyanatz, A., Cardenas, J. C., Gurven, M., Gwako, E., Henrich, N., Lesorogol, C., Marlowe, F., Tracer, D., & Ziker, J. (2010). Markets, religion, community size, and the evolution of fairness and punishment. *Science (New York, N.Y.)*, 327(5972), 1480–1484.

<https://doi.org/10.1126/science.1182238>

Hewstone, M., Rubin, M., & Willis, H. (2002). Intergroup bias. *Annual Review of Psychology*, 53, 575–604.

http://www.cairn.info/resume.php?ID_ARTICLE=RFG_170_0155%5Cnhttp://cat.inist.fr/?aModele=afficheN&cpsidt=3267136%5Cnhttp://ram.sagepub.com/content/18/4/3.short%5Cnhttp://onlinelibrary.wiley.com/doi/10.1111/j.1936-4490.2002.tb00667.x/abstract%5Cnhttp://scholar.goo

Jordan, J. J., Hoffman, M., Bloom, P., & Rand, D. G. (2016). Third-party punishment as a costly signal of trustworthiness. *Nature*, 530(7591), 473–476.

<https://doi.org/10.1038/nature16981>

Jordan, J. J., Mcauliffe, K., & Warneken, F. (2014). Development of in-group favoritism in children’s third-party punishment of selfishness. *Proceedings of the National Academy of Sciences*, 1–6. <https://doi.org/10.1073/pnas.1402280111>

Kosmidis, I. (2020). *brglm2: Bias Reduction in Generalized Linear Models*. R package

version 0.6.2. <https://cran.r-project.org/package=brglm2>

Kubota, J. T., Li, J., Bar-David, E., Banaji, M. R., & Phelps, E. A. (2013). The Price of Racial Bias. *Psychological Science*, 24(12), 2498–2504.

<https://doi.org/10.1177/0956797613496435>

Leach, C. W., van Zomeren, M., Zebel, S., Vliek, M. L. W., Pennekamp, S. F., Doosje, B., Ouwerkerk, J. W., & Spears, R. (2008). Group-level self-definition and self-investment: A hierarchical (multicomponent) model of in-group identification. *Journal of Personality and Social Psychology*, 95(1), 144–165.

<https://doi.org/10.1037/0022-3514.95.1.144>

Liddell, T. M., & Kruschke, J. K. (2014). Ostracism and fines in a public goods game with accidental contributions: the importance of punishment type. *Judgement and Decision Making*, 9(6), 523–547.

<http://econpapers.repec.org/RePEc:jdm:journl:v:9:y:2014:i:6:p:523-547>.

Martin, J. W., & Cushman, F. (2015). To Punish or to Leave: Distinct Cognitive Processes Underlie Partner Control and Partner Choice Behaviors. *PLOS ONE*, 10(4), e0125193. <https://doi.org/10.1371/journal.pone.0125193>

Martin, J. W., & Cushman, F. (2016). The Adaptive Logic of Moral Luck. *A Companion to Experimental Philosophy*, 190–202.

<https://doi.org/10.1002/9781118661666.ch12>

Matuschek, H., Kliegl, R., Vasishth, S., Baayen, H., & Bates, D. (2017). Balancing Type I error and power in linear mixed models. *Journal of Memory and Language*, 94, 305–315. <https://doi.org/10.1016/j.jml.2017.01.001>

McAuliffe, K., & Dunham, Y. (2016). Group bias in cooperative norm enforcement.

Philosophical Transactions of the Royal Society B: Biological Sciences, 371(1686), 20150073. <https://doi.org/10.1098/rstb.2015.0073>

McAuliffe, K., & Dunham, Y. (2017). Fairness overrides group bias in children's second-party punishment. *Journal of Experimental Psychology: General*, 146(4), 485–494.

<https://doi.org/10.1037/xge0000244>

McLeish, K. N., & Oxoby, R. J. (2011). Social interactions and the salience of social identity. *Journal of Economic Psychology*, 32(1), 172–178.

<https://doi.org/10.1016/j.joep.2010.11.003>

Mendoza, S. A., Lane, S. P., & Amodio, D. M. (2014). For Members Only. *Social Psychological and Personality Science*, 5(6), 662–670.

<https://doi.org/10.1177/1948550614527115>

Molenberghs, P., Gapp, J., Wang, B., Louis, W. R., & Decety, J. (2014). Increased Moral Sensitivity for Outgroup Perpetrators Harming Ingroup Members. *Cerebral Cortex*, 1–9. <https://doi.org/10.1093/cercor/bhu195>

Noe, R., & Hammerstein, P. (1994). Biological Markets : Supply and Demand Determine the Effect of Partner Choice in Cooperation , Mutualism and Mating\ . *Behavioral Ecology & Sociobiology*, 35(1), 1–11.

Pennycook, G., & Rand, D. G. (2020). Who falls for fake news? The roles of bullshit receptivity, overclaiming, familiarity, and analytic thinking. *Journal of Personality*, 88(2), 185–200. <https://doi.org/10.1111/jopy.12476>

Raihani, N. J., & Bshary, R. (2015). The reputation of punishers. *Trends in Ecology &*

- Evolution*, 1–6. <https://doi.org/10.1016/j.tree.2014.12.003>
- Raihani, N. J., & Bshary, R. (2019). Punishment: one tool, many uses. *Evolutionary Human Sciences*, 1, e12. <https://doi.org/10.1017/ehs.2019.12>
- Riek, B. M., Mania, E. W., & Gaertner, S. L. (2006). Intergroup threat and outgroup attitudes: A meta-analytic review. *Personality and Social Psychology Review*, 10(4), 336–353. https://doi.org/10.1207/s15327957pspr1004_4
- Sääksvuori, L., Mappes, T., & Puurtinen, M. (2011). Costly punishment prevails in inter group conflict. *Proceedings of the Royal Society B: Biological Sciences*, 278(1723), 3428–3436. <https://doi.org/10.1098/rspb.2011.0252>
- Sacco, D. F., Bernstein, M. J., Young, S. G., & Hugenberg, K. (2014). Reactions to social inclusion and ostracism as a function of perceived in-group similarity. *Group Dynamics*, 18(2), 129–137. <https://doi.org/10.1037/gdn0000002>
- Schiller, B., Baumgartner, T., & Knoch, D. (2014). Intergroup bias in third-party punishment stems from both ingroup favoritism and outgroup discrimination. *Evolution and Human Behavior*, 35(3), 169–175. <https://doi.org/10.1016/j.evolhumbehav.2013.12.006>
- Shinada, M., Yamagishi, T., & Ohmura, Y. (2004). False friends are worse than bitter enemies: “Altruistic” punishment of in-group members. *Evolution and Human Behavior*, 25(6), 379–393. <https://doi.org/10.1016/j.evolhumbehav.2004.08.001>
- Stagnaro, M. N., Dunham, Y., & Rand, D. G. (2018). Profit Versus Prejudice: Harnessing Self-Interest to Reduce In-Group Bias. *Social Psychological and Personality Science*, 9(1), 50–58. <https://doi.org/10.1177/1948550617699254>

Tajfel, H. (1970). Experiments in intergroup discrimination. *Scientific American*, 223(5), 96–102. <https://doi.org/10.1038/scientificamerican1170-96>

Tajfel, H., & Turner, J. (1979). An integrative theory of intergroup conflict. *The Social Psychology of Intergroup Relations*.
[https://www.google.com/books?hl=iw&lr=&id=l2R7DwAAQBAJ&oi=fnd&pg=PA56&dq=41.%09Tajfel,+H.,+%26+Turner,+J.+C.+\(1979\).+An+integrative+theory+of+intergroup+conflict.+In+W.+G.+Austin+%26+S.+Worchel+\(Eds.\),+The+social+psychology+of+intergroup+relations+&ots=Az3](https://www.google.com/books?hl=iw&lr=&id=l2R7DwAAQBAJ&oi=fnd&pg=PA56&dq=41.%09Tajfel,+H.,+%26+Turner,+J.+C.+(1979).+An+integrative+theory+of+intergroup+conflict.+In+W.+G.+Austin+%26+S.+Worchel+(Eds.),+The+social+psychology+of+intergroup+relations+&ots=Az3)

Trivers, R. L. (1971). The evolution of reciprocal altruism. *The Quarterly Review of Biology*, 46(1), 35–57. <https://doi.org/10.1086/406755>

Valenzuela, A., & Srivastava, J. (2012). Role of Information Asymmetry and Situational Salience in Reducing Intergroup Bias. *Personality and Social Psychology Bulletin*, 38(12), 1671–1683. <https://doi.org/10.1177/0146167212458327>

van der Lee, R., Ellemers, N., Scheepers, D., & Rutjens, B. T. (2017). In or out? How the perceived morality (vs. competence) of prospective group members affects acceptance and rejection. *European Journal of Social Psychology*, 47(6), 748–762. <https://doi.org/10.1002/ejsp.2269>

Van Prooijen, J. W. (2006). Retributive reactions to suspected offenders: The importance of social categorizations and guilt probability. *Personality and Social Psychology Bulletin*, 32(6), 715–726. <https://doi.org/10.1177/0146167205284964>

Watson-Jones, R. E., Whitehouse, H., & Legare, C. H. (2016). In-Group Ostracism Increases High-Fidelity Imitation in Early Childhood. *Psychological Science*, 27(1),

34–42. <https://doi.org/10.1177/0956797615607205>

Wirth, J. H., & Williams, K. D. (2009). “They don’t like our kind”: Consequences of being ostracized while possessing a group membership. *Group Processes and Intergroup Relations*, 12(1), 111–127. <https://doi.org/10.1177/1368430208098780>

Wu, J., Balliet, D., & Van Lange, P. A. M. (2016). Gossip Versus Punishment: The Efficiency of Reputation to Promote and Maintain Cooperation. *Scientific Reports*, 6(1), 23919. <https://doi.org/10.1038/srep23919>

Yamagishi, T. (1986). The provision of a sanctioning system as a public good. *Journal of Personality and Social Psychology*, 51(1), 110–116.
<https://doi.org/10.1037/0022-3514.51.1.110>

Yudkin, D. A., Rothmund, T., Twardawski, M., Thalla, N., & Van Bavel, J. J. (2016). Reflexive intergroup bias in third-party punishment. *Journal of Experimental Psychology: General*, 145(11), 1448–1459. <https://doi.org/10.1037/xge0000190>