Social Psychological and Personality Science
1-9
© The Author(s) 2017
Reprints and permission:
sagepub.com/journalsPermissions.nav
DOI: 10.1177/1948550617699254
journals.sagepub.com/home/spp
SAGE

Profit Versus Prejudice: Harnessing Self-Interest to Reduce In-Group Bias

Michael N. Stagnaro[1], Yarrow Dunham[1], and David G. Rand[1]

Abstract

We examine the possibility that self-interest, typically thought to undermine social welfare, might reduce in-group bias. We compared the dictator game (DG), where participants unilaterally divide money between themselves and a recipient, and the ultimatum game (UG), where the recipient can reject these offers. Unlike the DG, there is a self-interested motive for UG giving: If participants expect the rejection of unfair offers, they have a monetary incentive to be fair even to out-group members. Thus, we predicted substantial bias in the DG but little bias in the UG. We tested this hypothesis in two studies ($N = 3{,}546$) employing a 2 (in-group/out-group, based on abortion position) $\times$ 2 (DG/UG) design. We observed the predicted significant group by game interaction, such that the substantial in-group favoritism observed in the DG was almost entirely eliminated in the UG: Giving the recipient bargaining power reduced the premium offered to in-group members by 77.5%.

Keywords

intergroup relations, economic games, self-interest, in-group favoritism, discrimination

In-group bias is a central feature of human behavior with widespread negative consequences for both individuals and society (Fiske, 2002). These pressing social issues have generated a copious body of research on discrimination and bias reduction (Cikara & Van Bavel, 2014; Everett, Faber, & Crockett, 2015; Galinsky & Moskowitz, 2000; Milkman, Akinola, & Chugh, 2012; Pedersen, Walker, & Wise, 2005; Stephan, 1999; Wilder, 1986), most frequently via attempts to affect attitude change. However, these efforts have not been particularly successful (for review, see Paluck & Green, 2009). One reason could be because intergroup attitudes are only modest predictors of intergroup discrimination (Greenwald, Poehlman, Uhlmann, & Banaji, 2009; Oswald, Mitchell, Blanton, Jaccard, & Tetlock, 2013; Wicker, 1969). Here we suggest that sidestepping attitudes and examining the contexts within which actual acts of bias occur can be a fruitful strategy. Specifically, we focus on a special class of cases in which self-interest, a force typically thought to undermine social welfare, can actually reduce the behaviors associated with in-group bias.

A long-standing (and controversial) economic theory argues that bias can be costly not only for the recipient of discrimination but also for the discriminator (Becker, 1957). For example, an employer who offers an unreasonably low starting salary to a qualified candidate due to their opposing ideological beliefs may fail to recruit that candidate and so see their business suffer. We do not dispute that such discrimination occurs. However, we suggest that in some cases, this very same context will have quite a different outcome: Spurring some individuals to tacitly recognize that market forces and their own economic self-interest make bias costly, hence leading them to restrain or otherwise compensate for their prejudice. In this work, we present empirical evidence in support of this idea.

In-Group Bias in Monetary Division Tasks

There is a long precedent in social psychology and related fields of focusing on behaviors that favor the in-group and discriminate against the out-group, ranging from complex sociological recreations (Sherif, 1961; Zimbardo, 1971) to field experiments (Bertrand & Mullainathan, 2004; Gaertner & Dovidio, 1986; Paluck, 2011) and controlled laboratory paradigms (Dunham & Emory, 2014; Locksley, Ortiz, & Hepburn, 1980; Tajfel, Billig, Bundy, & Flament, 1971). Here we focus on work utilizing the latter methodology, specifically experiments using monetary divisions. A number of such studies have shown clear evidence for in-group bias, whereby people tend to allocate more money to members of their own group (for recent meta-analyses, see Balliet, Wu, & De Dreu, 2014; Lane, 2016).

Seminal work from Tajfel, Billig, Bundy, and Flament (1971), for example, found that participants randomized into "minimal" groups based on supposed preference for different

[1] Department of Psychology, Yale University, New Haven, CT, USA

Corresponding Author:
Michael N. Stagnaro, Department of Psychology, Yale University, 2 Hillhouse Avenue, New Haven, CT 06511, USA.
Email: michael.stagnaro@yale.edu

paintings consistently chose monetary divisions that either maximized the absolute earnings of the in-group or maximized the in-group's earnings relative to the out-group (even though doing so reduced the in-group's absolute earnings). To explain this bias, Tajfel et al. (1971) proposed social identity theory, whereby people elevate the status of their group in order to increase their own self-image. Another series of classic studies, summarized in Yamagishi, Jin, and Kiyonari (1999), showed boundary conditions of this minimal group effect whereby bias in allocations disappeared when the recipients did not know the allocator's group identity. The authors argued that minimal group bias was driven by an *exchange heuristic* whereby participants expect to be favorably treated by other members of their own group, rather than by an actual preference for the in-group (in contrast to social identity theory).

More recent work has examined in-group bias in the context of allocating money between the participant *themselves* and a member of the in-group versus out-group (i.e., the "dictator game" [DG]; Forsythe, Horowitz, Savin, & Sefton, 1994). In the DG, participants decide how to split a sum of money between themselves and a recipient; once the decision has been made, the recipient has no choice but to accept this split. This design adds an element not present in more classic allocation tasks, as the desire to benefit a member of the in-group versus out-group is now pitted against the desire to benefit oneself.

Research using minimal group paradigms shows that people give away more money when the recipient is a randomly assigned in-group member compared to an out-group member (Liebe & Tutic, 2010; McAuliffe & Dunham, 2016; Yamagishi & Mifune, 2008). A recent study by Everett, Faber, and Crockett (2015) demonstrated the robustness of this effect, finding that participants made more prosocial choices for in-group members regardless of whether the decision was public or private and whether it was framed in terms of gains or losses.

Further evidence of in-group bias in the DG has been demonstrated with more salient groups engaged in real-world conflict. For example, Fershtman and Gneezy (2001) found bias between different Israeli ethnic groups, which they argued to be the result of (mistaken) out-group ethnic stereotypes. Similarly, Whitt and Wilson (2007) showed offers reflecting in-group bias between Muslims, Croats, and Serbs in postwar Bosnia-Herzegovina. And several studies have demonstrated bias in DG offers among Americans based on political affiliation (Ben-Ner, McCall, Stephane, & Wang, 2009; Dunham, Arechar, & Rand, 2016; Fowler & Kam, 2007; Rand et al., 2009).

Self-Interested Motivations for Giving

While prior work provides ample evidence of in-group bias in the context of monetary divisions—both third-party allocations and DG giving—these are *unilateral* divisions in which subjects are free to enact whatever preferences they might have. Subjects are not required to take into account the receiver's preferences or actions, as the receiver must accept whatever offer is made. Thus, there is no self-interested motive to give, and the subject is free to act as they choose.

There are, however, other monetary division tasks in which the receiver is not just a passive recipient. The most widely used is the ultimatum game (UG; Güth, Schmittberger, & Schwarze, 1982). In the UG, as in the DG, an individual makes a decision of how to split a sum of money. However, here the recipient can accept the split (such that both individuals receive their respective amounts), or they can reject it (in which case neither player receives anything). Thus, unlike in the DG, the desire to avoid rejection creates a self-interested motivation for subjects to give money in the UG: If the receiver will reject unfair offers, it is payoff maximizing for subjects to be fair.

Supporting this logic, a large body of evidence suggests that fair offers in the UG are motivated largely by self-interest, with recipients often rejecting low offers and senders correctly predicting this fact (for reviews, see Camerer & Thaler, 1995; Thaler, 1988). For example, Suleiman (1996) and Handgraaf, Van Dijk, Vermunt, Wilke, and De Dreu (2008) considered a modified version of the UG where the receiver's power was varied, such that rejection destroyed only a specified fraction of both player's earnings. They found that as the fraction destroyed increased (i.e., as the recipient's power grew), senders increasingly made fair offers—showing an attention to the recipient's ability to reject and thus a strategic motive for behaving fairly.

In another study conducted over several countries (Israel, Japan, the United States, and Yugoslavia), researchers decreased the ability for recipients to reject offers by creating competition (Roth, Prasnikar, Okuno-Fujiwara, & Zamir, 1991). In this market-like game, nine recipients formed a market and bid on the lowest possible offer they would accept. A single sender then selected the best (lowest) demand from the group and sent them their proposed amount—resulting in offers much lower (10%) than those found in common UGs ($\sim$30%). Furthermore, the authors found that in all cases, the most common offer was the offer that was payoff maximizing given the behavior of the recipients—implicating strategic motives. Moreover, Abbink et al. (2000) had multiple senders compete for a single recipient and found the inverse effect such that senders offered over 60% on average, in the hopes of having their offer selected.

Finally, Wells and Rand (2013) designed a variant of the UG in which the receiver always received a constant sum (equal to the full endowment) irrespective of their choice, effectively removing any prosocial motive from the sender's decision—the only reason to give a nonzero offer is to avoid rejection. Critically, only the sender was aware of this feature, and thus the recipient would assume that the division was still relevant to their earnings (and would behave in the same way as a receiver in the standard UG). Senders were just as likely to make fair offers in this modified version as in the standard UG, despite the lack of prosocial motivation—suggesting that in the standard UG, strategic motives play a central role in fair offers.

The Present Work

Here, we investigated the power of self-interest to mitigate in-group bias by comparing favoritism in offers made in the DG

versus the UG. Little attention has been paid to in-group bias in UG offers to date—prior work on the impact of group membership on the UG has focused on *recipient* behavior, rather than proposer offers. One such example comes from Kubota, Li, Bar-David, Banaji, and Phelps (2013), who showed that recipients perceived low out-group offers as more hostile, and therefore rejected such offers more frequently compared to low in-group offers. Valenzuela and Srivastava (2012) found that recipients were more likely to accept both fair and unfair in-group offers compared to out-group offers. Mendoza, Lane, and Amodio (2014), conversely, found no group differences in recipients' acceptance of *overtly* fair or unfair offers but did find that recipients were more likely to reject *marginally* unfair offers from the in-group. Because these studies did not examine offers, it is unclear whether offers reflect tacit knowledge of these intergroup dynamics. One relevant study of proposers is Fershtman and Gneezy (2001), who found higher UG offers for members of groups that were seen as more insular and from "tough" honor-based societies, presumably because in such cases low offers could be perceived as an insult and spur rejection, but the link between these findings and intergroup bargaining remains unclear.

These findings motivate the following predictions: (1) in the DG, where recipients cannot penalize low offers, proposers matched with in-group recipients will give more than proposers matched with out-group members (i.e., in-group bias will be exhibited) and (2) proposers in the UG, tacitly understanding that low offers are at risk of rejection (regardless of the recipient's group membership), will make similar offers whether matched with in-group or out-group recipients—thus overriding in-group bias in order to avoid rejection.

Study 1

Method

Participants

U.S. residents ($N = 968$, 38.0% female, mean age 29.7 years) were recruited through Amazon Mechanical Turk (Horton, Rand, & Zeckhauser, 2011). Following standard practice in our lab for online studies, we aimed to recruit 250 participants per cell, to achieve a minimum of 200 participants per cell after excluding noncomprehenders. Due to the ease of collecting on Amazon Mechanical Turk and ongoing concerns about the replicability of psychological results, we also tested the robustness of our results by doubling the sample after the completion of the first collection (second collection: $N = 1{,}000$, 41.4% female, mean age 30.4 years). We focus on results for the full (doubled) sample but also report results broken down by collection.

Design

We used a 2 (in-group vs. out-group) × 2 (DG vs. UG) between-participants design. Abortion position was used as our group manipulation. Past work has shown that such political/ideological divisions can be highly polarizing and elicit similar levels of bias for both liberal and conservative individuals (Brandt, Reyna, Chambers, Crawford, & Wetherell, 2014; Wetherell, Brandt, & Reyna, 2013). Participants first provided demographics (age, race, gender, political, and abortion position: "pro-choice" or "pro-life"). They were then randomly assigned to read instructions for the DG or UG, followed by comprehension questions. Only comprehending participants (66% in the UG and 89% in the DG) were included in our analysis (imbalance in comprehension rates is addressed in Study 2). Next, participants were shown their partner's abortion position (randomized to be the same or opposite of the participant's) and told that the recipient was informed of their own position. Participants then chose how to split 40 cents with the partner. On the next screen, UG participants were asked to predict the minimum offer their recipient would have accepted ("minimal acceptable offer" [MAO]).

Compensation in Study 1 was determined as follows. After making their Role A decision (and predicting their partner's MAO in the UG condition), participants were taken to a new screen where they were told that they would play the same game again with a new partner, this time in Role B. Thus, each participant played the game twice (once in each role), and after all data had been collected, participants were matched based on abortion position and experimental condition and paid for both games they played. Because group condition was randomly assigned, half of people were told their partner was pro-choice and the other half pro-life; however, there were in fact more pro-choice participants ($N = 1{,}516$) than pro-life participants ($N = 452$). As a result, some subset of participants was not able to actually be matched with a partner who fit the description they were shown. In these cases, participants were sent a message stating that there had been an error in the matching algorithm and were paid the maximum bonus. Participants had no knowledge of the existence of the second decision while they were in Role A (the choices we actually analyzed), and thus there is no risk of this matching procedure affecting our results.

Results

As predicted, group bias was evident in the DG, but not in the UG. An analysis of variance (ANOVA) revealed significant main effects of game, $F(1, 1522) = 62.48$, $p < .0001$, with higher offers made in the UG ($M = .445$, standard deviation [SD] = .125) than the DG ($M = .369$, $SD = .225$), and group, $F(1, 1522) = 10.51$, $p = .001$, with higher offers made to in-group ($M = .419$, $SD = .178$) than out-group ($M = .383$, $SD = .204$) as well as the predicted Game × Group interaction, $F(1, 1522) = 5.63$, $p = .018$; effect size = 4.6 percentage points (Figure 1a).

Decomposing this interaction showed that while there was significant in-group favoritism in the DG, $t(865) = -3.59$, $p = .0004$ (in-group offers, $M = .396$, $SD = .215$; out-group offers, $M = .341$, $SD = .232$; effect size: 5.5 percentage points), this bias was almost entirely eliminated in the

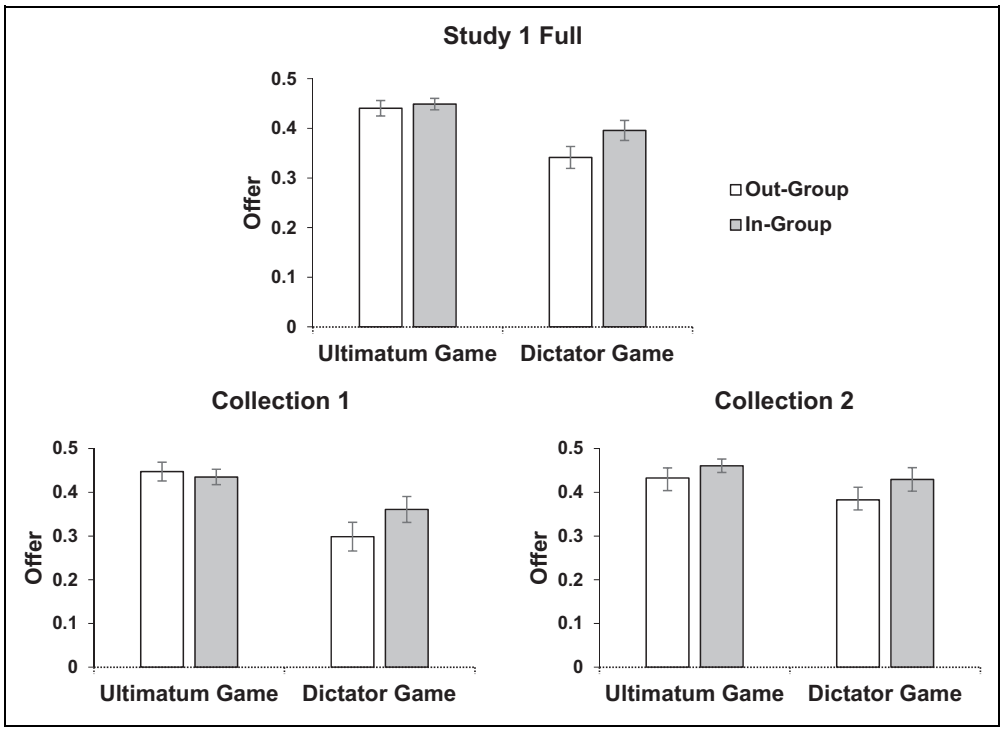

Figure 1. (a–c) Fraction of the endowment offered by participants in Study 1, based on game and partner's group. Error bars indicate 95% confident intervals.

UG, $t(654) = -0.86$, $p = .393$ (in-group offers, $M = .449$, $SD = .11$; out-group offers, $M = .440$, $SD = .14$; effect size = 0.8 percentage points).

Breaking this analysis down by collection (Figure 1b and c), we see a qualitatively similar pattern, with the difference between offers to in-group and out-group being bigger in the DG than the UG. However, the Game × Group interaction was only significant in Collection 1, $F(1, 750) = 6.88$, $p = .0089$, and not in Collection 2, $F(1, 771) = .52$, $p = .471$. At the same time, the three-way interaction between game, group, and collection was not significant, $F(1, 1522) = 2.11$, $p = .147$, indicating that the size of the Game × Group interaction did not differ significantly between collections (and that collapsing across collections was appropriate).

Evidence that this reduction in bias was driven by the desire to avoid rejection in the UG comes from two sources. First, we found a substantial positive correlation between participants' offers and their expectations about the minimum offer their recipient would accept, $r(654) = .38$, $p < .001$, as one would expect based on strategic reasoning. Second, expectations about the recipient's MAO did not significantly differ between in-group and out-group recipients, $t(652) = -0.967$, $p > .3$ (in-group, $M = .353$, $SD = .144$; out-group, $M = .342$, $SD = .151$; effect size = 1.1 percentage points; Figure 2). Thus, participants expected that they would need to make fair offers, regardless of the recipient's group, and (as is rational on the basis of this expectation) then actually made equivalent offers to in-group and out-group recipients.

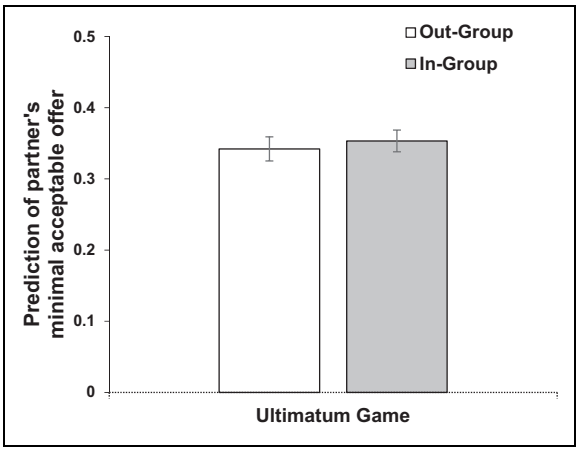

Figure 2. Study 1 participants' prediction of their partner's minimal acceptable offer in the ultimatum game. Error bars indicate 95% confident intervals.

Study 2

Study 2 differed from Study 1 in two ways. First, in Study 2, the recipient was described in more detail (beyond just abortion position), providing a more subtle manipulation of group affiliation. Second, the unbalanced comprehension rates between games were addressed; this imbalance occurred because there was a higher rate of comprehension check failure in the more complicated UG. To address this, after completing the UG or DG (depending on condition), participants read the instructions for the other game and answered the associated

comprehension questions. Only participants answering both sets of questions correctly were included in our analyses, resulting in a better balance in comprehension between games (UG: 51%, DG: 59%; these rates of comprehension are in line with what is typically observed in economic games run through Amazon Mechanical Turk [e.g., Rand, Greene, and Nowak, 2012]). Expecting a higher rate of comprehension failure compared to Study 1, we aimed to recruit 400 participants per cell for Study 2.

Method

Participants

We recruited $N = 1{,}578$ U.S. residents (43.5% female, mean age = 31.5 years) through Amazon Mechanical Turk. As before we used a 2 (in-group vs. out-group) × 2 (DG vs. UG) design.

Design

Group assignment. The procedure was identical to Study 1 through the first set of comprehension questions, after which participants were told to confirm their own demographic information before it was shared. Next, participants were told that their partner was 21–30 years old, English-speaking, male, Caucasian, and had a bachelor's degree; then, based on condition, was either (i) a Christian, pro-life, conservative or (ii) an atheist, pro-choice, liberal. As before, group affiliation was based on abortion position. Pro-life participants in the in-group condition and pro-choice participants in the out-group condition saw (i); pro-life participants in the out-group condition and pro-choice participants in the in-group condition saw (ii). These "richer" Player B descriptions were intended to minimize the threat of task demands by making it less obvious what group membership was being manipulated and to increase participants' credulity about the existence of an actual partner.

Game. The instructions for both the UG and DG were identical to those used in Study 1, except that participants had 100 cents to divide rather than 40 cents. After completing either the DG or UG, participants completed the comprehension questions for the other game.

For compensation in Study 2, we again had participants make Role B decisions after completing Role A. We avoided the matching issue experienced in Study 1 by modifying the procedure in two ways: When making their Role B decision (again, after making the key Role A decision), participants were (i) told that they might play the second game multiple times with different partners or might not actually play the second game at all and (ii) for this decision, they were only informed that their partner was "a worker on MTurk" (rather than giving specific information about the partner's abortion position, demographics, etc.)—thus avoiding the need to match on these characteristics for Role A players. We then identified the Role B decisions for some subset of participants who matched the partner description given to Role A players: 21- to 30-year-old English-speaking male Caucasians with a bachelor's degree who were either (i) Christian, pro-life, and conservative or (ii) atheist, pro-choice, and liberal. This subset was then used as the Role B players for all Role A players (i.e., they played the game many times, being paid for each, while the nonmatching people were not paid at all for the Role B game). Again, participants had no knowledge of the existence of the second decision or matching scheme while making the Role A choices analyzed in this article.

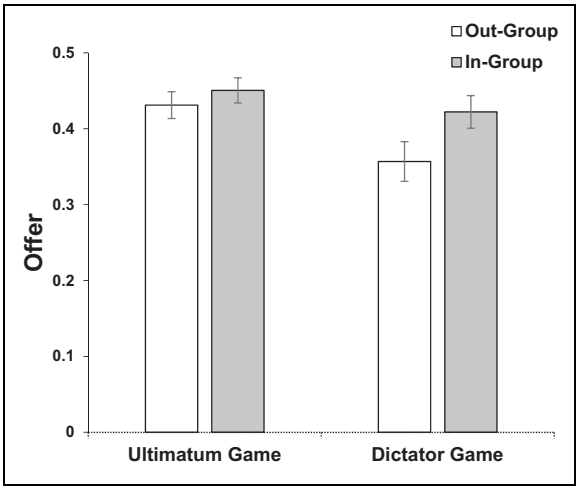

Figure 3. Fraction of the endowment offered by participants in Study 2, based on game and partner's group. Error bars indicate 95% confident intervals.

Results

Study 2 replicated the results of Study 1. An ANOVA revealed a significant main effect of game, $F(1, 871) = 20.64$, $p < .0001$, with higher offers made in the UG ($M = .441$, $SD = .12$) than the DG ($M = .387$, $SD = .196$) and group, $F(1, 871) = 14.05$, $p = .0002$, with higher offers made to in-group recipients ($M = .436$, $SD = .144$) than to out-group recipients ($M = .387$, $SD = .187$), as well as the predicted Game × Group interaction, $F(1, 871) = 4.16$, $p = .042$; effect size = 4.6 percentage points (Figure 3).

As in Study 1, decomposing the Game × Group interaction showed that while there was significant in-group favoritism in the DG, $t(486) = -3.715$, $p = .0002$ (in-group offers, $M = .422$, $SD = .163$; out-group offers, $M = .357$, $SD = .216$; effect size = 6.5 percentage points), this bias was mitigated in the UG, $t(382) = -1.579$, $p = .115$ (in-group offers, $M = .451$, $SD = .118$; out-group offers, $M = .431$, $SD = .122$; effect size = 1.9 percentage points). Furthermore, as in Study 1, participants' expectations about recipients' MAO were positively correlated with offers, $r(384) = .41$, $p < .001$, and did not significantly differ between recipients' group affiliation, $t(382) = 1.27$, $p > .2$ (expected demand from in-group, $M = .322$, $SD = .157$; from out-group, $M = .343$, $SD = .164$; effect size = 2.1 percentage points; Figure 4). Thus, in Study 2, as in Study 1, participants expected that they would

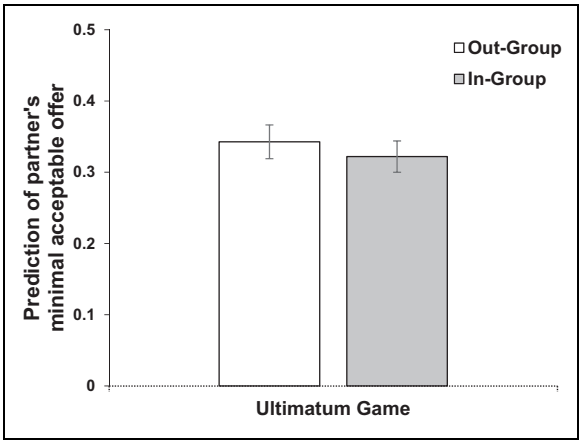

Figure 4. Study 2 participants' prediction of their partner's minimal acceptable offer in the ultimatum game. Error bars indicate 95% confident intervals.

need to make fair offers in order to avoid rejection, regardless of their partner's group affiliation.

Finally, we consider Study 2's presentation of information *beyond* just that which we manipulated for group identity (religion, political affiliation, and abortion position). This extra information (i.e., gender, race, age, language, and education) was added to create a more realistic partner description. However, adding this additional content could dilute our group manipulation effect by providing additional dimensions of similarity or difference between the participant and recipient beyond abortion position. For example, a participant who is pro-life and white in the out-group condition could be matched with a recipient who was pro-choice and white and is therefore out-group along the abortion dimension but in-group along the race dimension. To explore the possibility of dilution in such cases, we constrained our analysis to just those n = 88 participants who matched all the nonvarying features (gender, race, age, language, and education) of the presented recipient, and so who represent the cleanest test of our hypothesis. We continued to find the predicted Game × Group interaction, $F(1, 87) = 4.24$, $p = .043$; effect size = 11.7 percentage points, reflecting in-group favoritism in the DG, $t(48) = -2.368$, $p = .022$ (in-group offers, $M = .431$, $SD = .119$; out-group offers, $M = .3$, $SD = .232$; effect size = 13.1 percentage points), but not in the UG, $t(36) = 0.398$, $p = .693$ (in-group offers, $M = .45$, $SD = .111$; out-group offers, $M = .464$, $SD = .099$; effect size = 1.4 percentage points). Indeed, in this context, the raw effect sizes suggest an even more substantial effect of game on group bias, providing further support for our hypothesis.

Aggregated Results

Finally, given the similarity in design of both studies, we aggregated data from both to provide a higher powered test of our focal hypotheses. This analysis produced even starker results. In addition to main effects of game, $F(1, 2394) = 80.99$, $p < .0001$, and group, $F(1, 2394) = 26.48$, $p < .0001$, an ANOVA on the aggregated data found that the predicted Game × Group interaction was highly significant, $F(1, 2394) = 8.93$, $p = .0028$; effect size = 4.5 percentage points, such that while there was significant in-group favoritism in the DG, $t(1,353) = -4.95$, $p < .0001$ (in-group offers, $M = .405$ $SD = .199$; out-group offers, $M = .347$ $SD = .226$; effect size = 5.7 percentage points), this bias was mitigated in the UG, $t(1,038) = -1.63$, $p = .103$ (in-group offers, $M = .45$ $SD = .113$; out-group offers, $M = .437$ $SD = .134$; effect size = 1.3 percentage points). Thus, we find over the course of two studies and 3,546 participants that adding a strategic self-interest motive for fair offers (by giving the recipient bargaining power in the UG) reduced the premium offered to in-group members by 77.5%. Furthermore, as shown in the Supplemental Material, this effect was robust to including noncomprehenders, controlling for demographics, and constraining to participants who believed their recipient was real, and did not differ in size between pro-life and pro-choice participants.

Discussion

Here we have provided evidence that self-interest has the potential to override in-group bias based on a salient and highly charged real-world grouping (abortion stance). In the DG, where participants had the power to offer whatever they liked, we saw clear evidence of behavior favoring in-group members. In the UG, where the recipient could reject the offer, acting on such biases had the potential to severely reduce earnings. Participants anticipated this, as shown by their expectations of partner behavior, and made fair offers to both in-group and out-group participants.

Traditionally, self-interest is considered a negative force in intergroup relations. For example, an individual might give free reign to a preference for interacting with similar others, and even be willing to pay a cost to satisfy those preferences, resulting in what has been called "taste-based" discrimination (Becker, 1957). Although we do not deny that such discrimination can (and often does) occur, we suggest that in the right context, the costs it can impose serve as a disincentive. In particular, when strategic concerns are heightened, as they are in multilateral interactions where the parties must come to an agreement and failing to do so is both salient and costly (such as the UG), self-interest has the opportunity to mitigate biased behavior. Here, we provide one example of such a situation: We find that participants successfully withheld bias in the UG, making equally fair offers to both in-group and out-group recipients. When such strategic concerns are removed, as in the DG, the same pool of participants showed clear evidence of in-group favoritism in offers. It is also important to note that this strategic element, which is central for UG offers, does not exist for UG *responder* behavior, where bias has been observed previously (Kubota, Li, Bar-David, Banaji, & Phelps, 2013; Mendoza, Lane, & Amodio, 2014), because once the offer has been made, respondents can no longer influence offers. Indeed, regardless of what one is offered, rational self-interest never favors rejecting in the one-shot UG.

Thus, we observed a reduction in biased *behavior* in the UG when we introduced a self-interested reason to behave fairly, despite the *desire* to act in a biased way implied by the DG data. Importantly, although we provide evidence that incentivizing self-interest can mitigate in-group bias, we do not claim to have changed biased opinions. Indeed, given the relatively weak link between prejudicial attitudes and discriminatory behavior (Paluck & Green, 2009), our goal was to find a means to shift behavior without appealing to attitudes or directly attempting to foster attitude change. However, it is possible that over the long run, selfish incentives to act prosocially toward out-group members can actually reduce intergroup bias. This is because strategically induced changes in behavior can lead to substantial changes in internalized attitudes (Peysakhovich & Rand, 2016; Stagnaro, Arechar, & Rand, 2017), as it is adaptive to adjust one's default response (or "social heuristic"; Rand, 2016) to match behavior that is typically optimal (Bear, Kagan, & Rand, 2017; Bear & Rand, 2016).

By focusing on the nature of the intergroup interaction, and in particular whether it involves unilateral action or bilateral *coordination*, our work also highlights the need for continued focus on the specific structural features that characterize intergroup interactions. When control of outcomes is distributed, rather than residing wholly in one individual's hands, the costs associated with engaging in discrimination—and the extent to which it actually appears, at least in our data—are radically changed. This observation can be connected to classic theories of intergroup relations. For example, the need to, in essence, "reach a deal" that characterizes the UG may operate as a form of superordinate or common goal, which has been identified as a promising means of reducing prejudice in threat-based (Stephan & Stephan, 2000) as well as contact-based (Pettigrew & Tropp, 2005) approaches. However, in the present case, the goal motivating each individual is to reach a deal that maximizes their own *individual* profit. This differs from the standard approach in prejudice reduction, in which both individuals pursue the same goal, the fulfillment of which will benefit both individuals (e.g., to solve a puzzle together, earning a reward).

Similarly, some past research has focused on the role of status differences in producing prejudice and discrimination. For example, contact-based approaches highlight the need for contact to occur under conditions of equal status (Pettigrew & Tropp, 2005), and status differences themselves are robustly associated with greater prejudice (Bettencourt, Charlton, Dorr, & Hume, 2001). Although it is not typically conceptualized as such in the economic literature, the difference between DG and UG can be likened to a difference in status or power (Handgraaf, Van Dijk, Vermunt, Wilke, & De Dreu, 2008): In the DG, one player solely determines the outcome, while in the UG, the outcome is determined via an interaction of both players' behavior. As we discussed above, this interdependence highlights self-interested strategic concerns in a way that mitigates the free expression of prejudice.

Interestingly, recent developmental work (McAuliffe & Dunham, 2017) suggests that the pattern observed here might have deep roots in ontogeny: Children in randomly assigned novel social groups indicated a desire to preferentially interact with in-group members (through playmate preference) but did not show any evidence of biased offers in a modified, child-friendly UG (despite evidence that they do show such bias in the context of DGs, e.g., Benozio & Diesendruck, 2015). Thus, mitigating bias by invoking self-interest may not require extensive experience with bilateral bargaining.

In sum, we provide evidence that self-interest, typically seen as negatively impacting social outcomes including discrimination, can, at least in some circumstances, be harnessed for good: By placing our preferences for the group second to the goals of the self, strategic interactions may force us to be fair even to those whom we dislike. Although this logic will certainly not extend to all situations in which discrimination is present, our results suggest that strategic incentives may sometimes be a useful tool for those aiming to reduce bias.

Authors' Note

For all experiments, we have reported all measures, conditions, data exclusions, and how sample size was determined.

Declaration of Conflicting Interests

The author(s) declared no conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: The authors gratefully acknowledge funding from the Templeton World Charity Foundation (grant no. TWCF0209) and the Defense Advanced Research Projects Agency NGS2 program (grant no. D17AC00005).

Supplemental Material

The online data supplements are available at http://journals.sagepub.com/doi/suppl/10.1177/1948550617699254.

References

Abbink, K., Darzivj, R., Gilulaj, Z., Gorenj, H., Irlenbusche, B., Kerenj, A., . . . Zamirj, S. (2000). *The competitive ultimatum game: When competition matters, although it should not, and backward induction appears through the backdoor*. Retrieved from http://www.ma.huji.ac.il/∼zamir/papers/CugPap4.PDF

Balliet, D., Wu, J., & De Dreu, C. K. W. (2014). Ingroup favoritism in cooperation: A meta-analysis. *Psychological Bulletin*, *140*, 1556–1581. doi:10.1037/a0037737

Bear, A., Kagan, A., & Rand, D. G. (2017). Co-evolution of cooperation and cognition: The impact of imperfect deliberation and context-sensitive intuition. *Proceedings of the Royal Society B: Biological Sciences*, *284*, pii: 20162326. doi:10.1098/rspb.2016.2326

Bear, A., & Rand, D. G. (2016). Intuition, deliberation, and the evolution of cooperation. *Proceedings of the National Academy of Sciences*, *113*, 936–941. doi:10.1073/pnas.1517780113

Becker, G. S. (1957). *The economics of discrimination*. Chicago, IL: University of Chicago Press.

Ben-Ner, A., McCall, B. P., Stephane, M., & Wang, H. (2009). Identity and in-group/out-group differentiation in work and giving behaviors: Experimental evidence. *Journal of Economic Behavior & Organization*, *72*, 153–170.

Benozio, A., & Diesendruck, G. (2015). Parochialism in preschool boys' resource allocation. *Evolution and Human Behavior*, 1–9. doi:10.1016/j.evolhumbehav.2014.12.002

Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *The American Economic Review*, *94*, 991–1013.

Bettencourt, B., Charlton, K., Dorr, N., & Hume, D. L. (2001). Status differences and in-group bias: A meta-analytic examination of the effects of status stability, status legitimacy, and group permeability. *Psychological Bulletin*, *127*, 520.

Brandt, M. J., Reyna, C., Chambers, J. R., Crawford, J. T., & Wetherell, G. (2014). The ideological-conflict hypothesis intolerance among both liberals and conservatives. *Current Directions in Psychological Science*, *23*, 27–34.

Camerer, C., & Thaler, R. H. (1995). Anomalies: Ultimatums, dictators and manners. *The Journal of Economic Perspectives*, *9*, 209–219.

Cikara, M., & Van Bavel, J. J. (2014). The neuroscience of intergroup relations an integrative review. *Perspectives on Psychological Science 9*245–274

Dunham, Y., Arechar, A. A., & Rand, D. G. (2016). The temporal dynamics of intra-party bias in U.S. electoral politics. Retrieved from https://ssrn.com/abstract=2846915

Dunham, Y., & Emory, J. (2014). Of affect and ambiguity: The emergence of preference for arbitrary ingroups. *Journal of Social Issues*, *70*, 81–98. doi:10.1111/josi.12048

Everett, J. A., Faber, N. S., & Crockett, M. J. (2015). The influence of social preferences and reputational concerns on intergroup prosocial behavior in gains and losses contexts. *Royal Society Open Science*, *2*, 150546. doi:10.1098/rsos.150546

Fershtman, C., & Gneezy, U. (2001). Discrimination in a segmented society: An experimental approach. *Quarterly Journal of Economics*, *116*, 351–377.

Fiske, S. T. (2002). What we know now about bias and intergroup conflict, the problem of the century. *Current Directions in Psychological Science*, *11*, 123–128.

Forsythe, R., Horowitz, J. L., Savin, N. E., & Sefton, M. (1994). Fairness in simple bargaining experiments. *Games and Economic Behavior*, *6*, 347–369.

Fowler, J. H., & Kam, C. D. (2007). Beyond the self: Social identity, altruism, and political participation. *The Journal of Politics*, *69*, 813–827.

Gaertner, S. L., & Dovidio, J. F. (1986). *The aversive form of racism*. Orlando, FL: Academic Press, 61–89.

Galinsky, A. D., & Moskowitz, G. B. (2000). Perspective-taking: Decreasing stereotype expression, stereotype accessibility, and in-group favoritism. *Journal of Personality and Social Psychology*, *78*, 708.

Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., & Banaji, M. R. (2009). Understanding and using the implicit association test: III. Meta-analysis of predictive validity. *Journal of Personality and Social Psychology*, *97*, 17.

Güth, W., Schmittberger, R., & Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, *3*, 367–388.

Handgraaf, M. J., Van Dijk, E., Vermunt, R. C., Wilke, H. A., & De Dreu, C. K. (2008). Less power or powerless? Egocentric empathy gaps and the irony of having little versus no power in social decision making. *Journal of Personality and Social Psychology*, *95*, 1136.

Horton, J. J., Rand, D. G., & Zeckhauser, R. J. (2011). The online laboratory: Conducting experiments in a real labor market. *Experimental Economics*, *14*, 399–425.

Kubota, J. T., Li, J., Bar-David, E., Banaji, M. R., & Phelps, E. A. (2013). The price of racial bias intergroup negotiations in the ultimatum game. *Psychological Science*, *24*, 2498–2504.

Lane, T. (2016). Discrimination in the laboratory: A meta-analysis of economics experiments. *European Economic Review*, *90*, 375–402. doi:10.1016/j.euroecorev.2015.11.011

Liebe, U., & Tutic, A. (2010). Status groups and altruistic behavior in dictator games. *Rationality and Society*, *22*, 353–380. doi:10.1177/1043463110366232

Locksley, A., Ortiz, V., & Hepburn, C. (1980). Social categorization and discriminatory behavior: Extinguishing the minimal intergroup discrimination effect. *Journal of Personality and Social Psychology*, *39*, 773.

McAuliffe, K., & Dunham, Y. (2016). Group bias in cooperative norm enforcement. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *371*, 20150073. doi: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*, 485–494. doi:10.1037/xge0000244

Mendoza, S. A., Lane, S. P., & Amodio, D. M. (2014). For members only ingroup punishment of fairness norm violations in the ultimatum game. *Social Psychological and Personality Science*, *5*, 662–670.

Milkman, K. L., Akinola, M., & Chugh, D. (2012). Temporal distance and discrimination an audit study in academia. *Psychological Science*, *23*, 710–717.

Oswald, F. L., Mitchell, G., Blanton, H., Jaccard, J., & Tetlock, P. E. (2013). Predicting ethnic and racial discrimination: A meta-analysis of IAT criterion studies. *Journal of Personality and Social Psychology*, *105*, 171–192.

Paluck, E. L. (2011). Peer pressure against prejudice: A high school field experiment examining social network change. *Journal of Experimental Social Psychology*, *47*, 350–358.

Paluck, E. L., & Green, D. P. (2009). Prejudice reduction: What works? A review and assessment of research and practice. *Annual Review of Psychology*, *60*, 339–367.

Pedersen, A., Walker, I., & Wise, M. (2005). "Talk does not cook rice": Beyond antiracism rhetoric to strategies for social action. *Australian Psychologist*, *40*, 20–31.

Pettigrew, T. F., & Tropp, L. R. (2005). Allport's intergroup contact hypothesis: Its history and influence. *On the Nature of Prejudice*, *50*, 262–277.

Peysakhovich, A., & Rand, D. G. (2016). Habits of virtue: Creating norms of cooperation and defection in the laboratory. *Management Science*, *62*, 631–647.

Rand, D. G. (2016). Cooperation, fast and slow: Meta-analytic evidence for a theory of social heuristics and self-interested deliberation. *Psychological Science*, *27*, 1192–1206. doi:10.1177/0956797616654455

Rand, D. G., Greene, J. D., & Nowak, M. A. (2012). Spontaneous giving and calculated greed. *Nature*, *489*, 427–430. doi:10.1038/nature11467

Rand, D. G., Pfeiffer, T., Dreber, A., Sheketoff, R. W., Wernerfelt, N. C., & Benkler, Y. (2009). Dynamic remodeling of in-group bias during the 2008 presidential election. *Proceedings of the National Academy of Sciences of the United States of America*, *106*, 6187–6191.

Roth, A. E., Prasnikar, V., Okuno-Fujiwara, M., & Zamir, S. (1991). Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study. *The American Economic Review*, *81*, 1068–1095.

Sherif, M. (1961). *Intergroup conflict and cooperation: The Robbers Cave experiment* (Vol. *10*, pp. 150–198). Norman, OK: University Book Exchange.

Stagnaro, M. N., Arechar, A. A., & Rand, D. G. (2017). From good institutions to generous citizens: Top-down incentives to cooperate promote subsequent prosociality but not norm enforcement. *Cognition*. Advance online publication. doi:10.1016/j.cognition.2017.01.017

Stephan, W. G. (1999). *Reducing prejudice and stereotyping in schools*. New York, NY: Teachers College Press.

Stephan, W. G., & Stephan, C. W. (2000). An integrated threat theory of prejudice. In S. Oskamp (Ed.), *Reducing prejudice and discrimination* (pp. 23–45). New Jersey: Psychology Press.

Suleiman, R. (1996). Expectations and fairness in a modified ultimatum game. *Journal of Economic Psychology*, *17*, 531–554.

Tajfel, H., Billig, M. G., Bundy, R. P., & Flament, C. (1971). Social categorization and intergroup behavior. *European Journal of Social Psychology*, *1*, 149–178. doi:10.1002/ejsp.2420010202

Thaler, R. H. (1988). Anomalies: The ultimatum game. *Journal of Economic Perspectives*, *2*, 195–206.

Valenzuela, A., & Srivastava, J. (2012). Role of information asymmetry and situational salience in reducing intergroup bias: The case of Ultimatum Games *Personality and Social Psychology Bulletin*, *38*1671–1683

Wells, J. S., & Rand, D. G. (2013). *Strategic self-interest can explain seemingly "fair" offers in the Ultimatum Game*. Retrieved from SSRN 2136707

Wetherell, G. A., Brandt, M. J., & Reyna, C. (2013). Discrimination across the ideological divide the role of value violations and abstract values in discrimination by liberals and conservatives. *Social Psychological and Personality Science*, *4*, 658–667.

Whitt, S., & Wilson, R. K. (2007). The dictator game, fairness and ethnicity in postwar Bosnia. *American Journal of Political Science*, *51*, 655–668. doi:10.1111/j.1540-5907.2007.00273.x

Wicker, A. W. (1969). Attitudes versus actions: The relationship of verbal and overt behavioral responses to attitude objects. *Journal of Social Issues*, *25*, 41–78.

Wilder, D. A. (1986). Social categorization: Implications for creation and reduction of intergroup bias. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. *19*, pp. 239–355). New York, NY: Academic Press.

Yamagishi, T., Jin, N., & Kiyonari, T. (1999). Bounded generalized reciprocity: Ingroup boasting and ingroup favoritism. *Advances in Group Process*, *16*, 161–197.

Yamagishi, T., & Mifune, N. (2008). Does shared group membership promote altruism? Fear, greed, and reputation. *Rationality and Society*, *20*, 5–30.

Zimbardo, P. G. (1971). *The power and pathology of imprisonment* (Congressional Record; Serial No. 15, 1971–10–25). Washington, DC: U.S. Government Printing Office.

Author Biographies

Michael N. Stagnaro is a fourth-year PhD student in psychology. His work combines social, cognitive, and evolutionary psychology to address questions at the intersection of morality and religion and their effects on cooperation and defection as well as the role institutions play in the development of these topics.

Yarrow Dunham is an assistant professor of psychology and cognitive science and the director of the SCD Lab. His research focuses on studying how knowledge of social groups is acquired, both in cognitively mature adults and in the developing children. His research combines a range of experimental and cross-cultural methodologies.

David G. Rand is an associate professor of psychology, economics, and management at Yale University and the director of Yale University's Human Cooperation Laboratory. His research combines a range of theoretical and experimental methods in an effort to explain the high levels of cooperation that typify human societies.

Handling Editor: Joseph Simmons