

Deliberation erodes cooperative behavior – even towards competitive out-groups, even when using a control condition, and even when eliminating selection bias

Jim A.C. Everett^{1,2} ✉, Zach Ingbreetsen¹, Fiery Cushman¹, Mina Cikara¹

¹Department of Psychology, Harvard University, Cambridge, MA, USA.,

²Department of Experimental Psychology, Oxford University, Oxford, UK.,

✉ Corresponding author: jim.ac.everett@gmail.com

By many accounts cooperation appears to be a default strategy in social interaction. There are, however, several documented instances in which reflexive responding favors aggressive behaviors: for example, interactions with out-group members. We conduct a rigorous test of potential boundary conditions of intuitive prosociality by looking at whether intuition favors cooperation even towards competitive out-group members, and even in losses frames. Moreover, we address three major methodological limitations of previous research in this area: a lack of an unconstrained control condition; non-compliance with time manipulations leading to high rates of exclusions and thus a selection bias; and non-comprehension of the structure of the game. Even after eliminating participant selection bias and non-comprehension, we find that deliberation decreases cooperation: even in competitive contexts towards out-groups and even in a losses frame, though the differences in cooperation between groups was consistent across conditions. People may be intuitive cooperators, but they are not intuitively impartial.

A recent body of research suggests that, far from requiring effortful control, behaving prosocially arises from “processes that are intuitive, reflexive, and even automatic” (Zaki & Mitchell, 2013, p. 466). Several studies find that people tend to make prosocial decisions in economic games more quickly than selfish ones, and time-pressure increases the incidence of prosocial behavior (for an overview and meta-analysis see Rand, 2016). Time delay reduces helping in a ‘dropped-glove’ field study (Artavia-Mora, Bedi, & Rieger, 2017) and even risking one’s own life to save another seems driven primarily by intuitive processes (Rand & Epstein, 2014). To explain this, the Social Heuristics Hypothesis (SHH: Rand, 2016; Rand et al., 2014) posits that the social strategies typically successful in daily life (e.g., cooperation) become automatized as default responses, and that deliberation can override these defaults to modify behavior. Indeed, a formal analysis of evolutionary dynamics indicates that

deliberation can only serve to undermine costly cooperation and not promote it (Bear, Kagan, & Rand, 2017; Bear & Rand, 2016). While meta-analytic work has provided strong support for the claim that manipulating reliance on intuition through time pressure encourages prosociality (Rand, 2016; Rand et al., 2014), a recent Registered Replication Report (RRR) by Bouwmeester and colleagues (2017) finds only mixed support. Our aim is to consider two challenges to the idea that intuition favors cooperation—one methodological and one theoretical—and then to provide new experimental evidence.

Boundary conditions on intuitive cooperation: Intergroup bias and decision framing

One of the most enduring findings in social psychology is intergroup bias: the powerful tendency to evaluate and treat in-group members more favorably than out-group members (Hewstone, Rubin, & Willis, 2002). How does intuitive

cooperation play out in intergroup contexts? The SSH posits that cooperative strategies that are typically advantageous in daily life with repeated trustworthy interaction partners become internalized as intuitions and get overgeneralized to less typical settings. Given this, we might reason that social heuristics, precisely because they are overgeneralized, will apply even to out-group members and it is deliberation that ‘corrects’ this process and causes increased intergroup bias. Supporting this, for example, is an experiment employing a Public Goods Game with minimal groups; conceptual priming of intuition leads to equivalent contributions to both in-group and out-group members, whereas priming reflection leads to pronounced in-group favoritism (Ma, Liu, Rand, Heatherton, & Han, 2015).

Yet the idea that intuition uniformly favors cooperation even with out-group members is at odds with much past evidence suggesting that the preference for ‘us’ over ‘them’ is at least partially rooted in implicit, unconscious, or automatic processes. Preference for in-group members appears early in development (Dunham, Baron, & Banaji, 2008), and group memberships appear to generate bias at the earliest stages of perceptual and emotional processes (for a review see Cikara & Van Bavel, 2014). Intergroup bias often occurs outside the realm of conscious awareness and, for this reason, can be remarkably difficult to control or change permanently (Greenwald, McGhee, & Schwartz, 1998). For example, a recent test of 9 interventions that reduce implicit prejudice in the IAT found that none lasted more than one day (Lai et al., 2016). There is also evidence for intuitive aggression towards out-groups - for example, in the Shooter task where participants are required to “shoot” armed targets and to “not shoot” unarmed targets (Correll, Park, Judd, & Wittenbrink, 2002). Participants under time pressure are more likely to mistakenly shoot other-race targets than same-race targets (Correll et al., 2002) and are more likely to mistakenly shoot minimal out-group members than in-group members (Miller, Zielaskowski, & Plant, 2012).

Looking specifically at prosocial behaviour, several convergent findings

indicate that intuition encourages in-group favoritism rather than impartial cooperation. Experiments using the Intergroup Prisoner’s Dilemma–Maximizing Differences Game (Halevy, Bornstein, & Sagiv, 2008) find that parochial altruism - contributions to a pool that both benefits the in-group and simultaneously hurts the out-group - emerges especially among individuals who were cognitively taxed by completing a Stroop interference task (Dreu, Dussel, & Velden, 2015). Another series of experiments using a Prisoner’s Dilemma with real political groups (Obama vs. Romney supporters) and a time pressure manipulation demonstrate that intuition increases cooperation to both in-group and out-group members, but that the difference between groups is maintained relative to an enforced time delay condition (Rand, Newman, & Wurzbacher, 2015). Therefore, while at least one study indicates that intuition reduces in-group favoritism (Ma et al., 2015), two others indicate that intuition increases in-group favoritism (De Dreu et al., 2015; Ten Velden, Daughters, & De Dreu, 2017), and a fourth study indicates that intuition increases prosocial behavior but neither increases nor decreases in-group favoritism (Rand et al., 2015).

One possibility for this inconsistency is that existing studies using time pressure have typically contrasted a manipulation intended to promote intuition against a manipulation intended to promote reflection, but do not assess these manipulations against an unconstrained control condition (e.g. Capraro, Jordan, & Rand, 2014; Cone & Rand, 2014; Rand et al., 2014; Verkoeijen & Bouwmeester, 2014). Such a design cannot distinguish evidence that time pressure increases cooperation from evidence that time constraint reduces cooperation; as such, discrepancies across the intergroup studies might reflect the relative efficacy of the intuitive versus deliberative manipulations included in each.

Our second theoretical contribution is to explore the consequences of framing a social dilemma in terms of losses versus gains. Motivated in part by prior research on active hostility in intergroup contexts (e.g. Correll et al., 2002; Sherif, 1966), we speculated that defection against out-

group members might be the intuitive response especially in a social dilemma framed in terms of losses (i.e., defection imposes a cost) rather than gains (i.e., defection withholds a benefit). Prior research on framing in social dilemmas is inconsistent, and has varyingly shown loss frames to increase cooperation (e.g. Experiment 3 in Komorita & Carnevale, 1992), reduce cooperation (e.g. Brewer & Kramer, 1986), or have no effect (e.g. de Heus, Hoogervorst, & Dijk, 2010). This may occur because decision frames have divergent effects based on an individual's prior motives, such that prosocial people become more cooperative, and 'individualists' less cooperative (De Dreu & McCusker, 1997). Because people's prior motives are different with regards to in-group and out-group members, we might therefore expect loss frames to encourage defection towards out-groups (because without a prior motive to be cooperative, loss frames make people more selfish) and cooperation towards in-groups (because loss frames enhance the existing motive to help the in-group member).

Methodological Challenges for Intuitive Cooperation

In addition to these theoretical concerns, recent research has challenged intuitive cooperation findings on methodological grounds (Tinghög et al., 2013). This critique focuses especially on the use of time pressure versus time delay to manipulate the balance of automatic versus controlled inputs into cooperation decisions, specifically citing high levels of participant exclusions and selection bias (Bouwmeester et al., 2017; Tinghög et al., 2013). The original studies excluded participants who failed to make their decisions within the response window. These exclusion rates are typically very high and, more importantly, asymmetric across conditions. For example, in Rand, Greene, and Nowak (2012), 48% of participants failed to make their decisions under time pressure in Study 6, and 46% failed in Study 7 relative to 19% and 10% in the time delay conditions, respectively. While by some accounts these effects hold when including non-compliant participants (Rand et al., 2014; Rand, Greene, & Nowak, 2013), these exclusion practices introduce the possibility of

selection bias. Eliminating participants who are too slow from the time pressure condition and not from the time delay condition disrupts random assignment to condition. Therefore, the observed difference in cooperation could be driven by systematic differences between the participants rather than the manipulation. For example, Tinghög et al. (2013) could not successfully replicate Rand et al.'s (2012) results without exclusions and therefore conclude the original findings were "an artefact of excluding the about 50% of subjects who failed to respond on time" (p.427). Consistent with this possibility, a recent pre-registered multi-site replication study (Bouwmeester et al., 2017) reported that two-thirds of participants failed to make decisions within the allotted time and that the effect of time pressure on cooperation was only present when excluding such individuals. Of course, this data is also consistent with the possibility that individuals who fail to conform to the time-pressure treatment therefore show no effect of that treatment.

A similar problem arises with the use of comprehension checks, where large numbers of participants fail to correctly answer comprehension questions about the structure of the game after they have played it. Across the studies reported by Rand et al. (2012), comprehension checks were implemented after the game had been played to avoid suggesting a deliberative mindset to participants prior to decision making. And indeed, as Rand and colleagues report in a supplementary study, participants who complete comprehension questions before making their decision choose to contribute significantly less than those who complete the comprehension questions afterward (Rand et al., 2012). However, in practice this means that participants may play the game while not understanding it. For example, 32% of participants in Study 1 of Rand et al. (2015) failed one or both comprehension checks and yet were included in the final analysis. While the effects of time pressure were robust to controlling for comprehension, this remains a potentially problematic aspect of the dominant methodological design.

Given these concerns (and the lack of the control condition raised in the previous section), we designed a procedure

that accomplishes three key methodological goals. First, we drastically reduced exclusion rates due to the response window, achieving an exclusion rate of just 2% of participants enrolled in the study. We accomplished this by providing participants with extensive comprehension training prior to the task employing alternative payoff matrices. This prepared them to quickly absorb and respond to information presented briefly in the main task. Second, we equated the exclusion procedure across the time pressure and time delay conditions. We accomplished this by subjecting participants in the time delay condition to an additional, subsequent time-pressure trial, and then excluding their data if they failed to meet the response deadline on this subsequent trial. Thus, our exclusion rates were both very low and comparable across conditions. Finally, we included an unconstrained control condition to test whether time delay, time pressure, or both are significantly different from participants' default responses.

Method

Pre-Registration

Our design, hypotheses, and analysis plan were all pre-registered at the Open Science Framework (<https://osf.io/v82gc/>). We report all measures, manipulations, and exclusions in this study. All data, analysis code, and experiment materials are available for download at: <https://osf.io/pz4he/>.

Participants

1316 American participants completed the experiment online using Amazon Mechanical Turk and received \$0.70 for their time, with all participants being paid the same additional bonus of \$0.70 (the maximum achievable in the task). In accordance with the pre-registration we planned, participants were excluded if they completed the survey more than once ($N = 158$); failed simple comprehension checks regarding which group they had been assigned to ($N = 0$); and/or failed to complete the main PD trial in the time pressure condition ($N = 9$) or failed to complete the time pressure bonus round in

the time delay condition ($N = 9$). Therefore, the final sample consisted of 1140 participants (557 female). For these main effects we had around 380 participants per condition for the three time-instruction conditions, and 570 per condition for the gains/loss and in-group/out-group conditions.

In order to assess our statistical power to detect the elimination of intuitive cooperation among outgroup targets, in a loss frame, or when these features are combined, we conducted three power analyses using Fisher's exact test in G*Power (Faul, Erdfelder, Buchner, & Lang, 2009). For the purposes of simplicity we excluded our baseline control condition from this simulation, focusing just on the effect of time pressure vs. delay. Assuming that the true effect is 50% cooperation under time delay and 70% cooperation under time pressure, setting $\alpha = .05$, and applying the actual sample sizes we obtained following participant exclusions, we conducted three analyses corresponding to those reported in the results section. First, for the effects of time pressure vs. time delay restricted to the outgroup conditions, we had power of .98. Second, for the the effects of time pressure vs. time delay restricted to the losses frame conditions, we had power of .97. Third, for the effects of time pressure vs. delay restricted to the outgroup and loss frame condition, we had power of .77.

Design

We used a fully between-subjects design where participants completed a Prisoner's Dilemma (PD; Axelrod, 1980; Rapoport & Chammah, 1965) and we manipulated (1) whether participants were encouraged to make their decision intuitively, through deliberation, or in the absence of either instruction; (2) the group membership of the other player; and (3) whether the PD was framed in terms of gains or losses. This study had a fully between-subjects design with 12 conditions: 2 (group membership of other player: in-group vs. out-group) x 2 (frame: gains vs. losses) x 3 (instructions: time delay vs. time pressure vs. unconstrained control).

Upon beginning the study, participants were told that they would be randomly

assigned to one of two teams (Eagles or Rattlers) based on their responses to five short personality questions, and that these teams were competing with other in a problem solving task (Cikara, Bruneau, Van Bavel & Saxe, 2014). Upon being assigned to a team (always the Eagles), participants were told that before completing the competitive problem-solving task they would first engage in a short and entirely separate game (the PD) with another participant, where this other participant could either be from their team (in-group), or the other team (out-group). We employed novel teams to control for the effects of stereotypes and familiarity, but made them competitive to simulate the functional relations between groups in conflict (Chang, Krosch, & Cikara, 2016).

Participants were then given comprehensive instructions and extensive training on the structure of the PD to ensure that they understood the different possible outcomes of the game (see materials on OSF). All participants were required to manually enter the different outcomes that each player would receive depending on the choices of the other player in an example matrix, and participants had to enter these correctly before being able to move forward in the study. Therefore, and in contrast to most previous studies, we did not have to exclude any participants based on failing to understand the structure of the game¹.

The PD was framed in one of two different ways (see Figure 1). In the first, standard, version of the game – the gains frame – participants learned that they could earn an additional bonus of up to \$0.70 depending on the decisions that both they and the other player made in the game. For example, if both participants chose to cooperate, they would both gain \$0.50 and so would end the game with a bonus of \$0.50. In the the losses frame participants were told they had been given a bonus of \$0.80 and that they would lose a certain amount (at least \$0.10) based on the decision they made and the decision the other player made. For example, if both participants chose to cooperate, they would both lose \$0.30 and would end the game with a bonus of \$0.50. Therefore, in both versions of the game, participant's choice to cooperate or defect would yield the same eventual bonus – but whether this

was framed as gaining or losing money differed.

To encourage participants to make their decision intuitively or through deliberation, we drew on previous research and had participants make their decision under time pressure or time delay (Cone & Rand, 2014; Rand, Brescoll, Everett, Capraro, & Barcelo, 2016; Rand et al., 2012). Reducing the amount of time participants have to make their decision decreases the potential for deliberation to outweigh intuition, thus leading to more intuitive decisions (Wright, 1974). In this study, we encouraged participants to rely on intuition by requiring them to make their decision in less than 15 seconds (time pressure), encouraged to think deliberatively by requiring them wait for at least 15 seconds before deciding (time delay), or did not specify how long to take before making a decision (control). This specific time to induce intuition vs. deliberation (15 seconds) was determined after running a pilot study (N = 102) in which, in the absence of any instructions to be fast or slow, participants took a median time of 17 seconds. As planned, participants in the intuition condition who did not make a decision within the allotted time were excluded from data analysis. Furthermore, in order to ensure that participants in the time pressure and time delay conditions did not systematically differ in terms of their ability to answer in a short time-frame, participants in the time delay condition completed a bonus round of the PD (with an alternative matrix in which all values were reduced by 0.05) under time pressure after completing the main task. Participants in the time delay condition who failed to make a decision in the allotted time in this bonus round were also excluded from data analysis. In the main task, participants in the time delay condition took a median time of 23 seconds, and participants in the intuition condition took a median time of 9 seconds. Response times for the control condition were not recorded due to a technical error.

Results

In order to assess the effects of group membership, gains/losses frame, and time manipulation on the likelihood that participants made cooperative decisions in

the PD we conducted a series of logistic regressions in R. We began by fitting the full summative model (including all predictors and their interactions) and then used likelihood ratio tests to compare its fit with various nested models. We started by dropping the interaction terms and found that none of these improved the model fit. However, dropping any of the main effects (group; gains/losses frame; time manipulation) did result in significantly worse fit, indicating that all three predictors were necessary for analysis. We therefore report the results of this best-fitting model, which includes additive effects of each predictor, but no interactions and examine simple effects using least-square (adjusted) means with the Tukey method to correct for multiple comparisons.

This model showed main effects of all three predictors (see Table 1 and Figure 2). There was a significant effect of decision frame ($B = -.42$, $SE = .12$, $Z = -3.52$, $p < .001$) such that participants were more likely to cooperate under a gains frame (61%) than a losses frame (51%), and a significant effect of group membership ($B = -.30$, $SE = .12$, $Z = -2.51$, $p = .01$) such that participants were more likely to cooperate with an in-group (60%) than an out-group member (53%). Finally, participants were more likely to cooperate under time pressure (63%) than under time delay (49%) ($B = .58$, $SE = 0.15$, $Z = 3.81$, $p = .004$), and marginally more likely to cooperate when given no instructions (57%) than under time delay ($B = .31$, $SE = 0.15$, $Z = 2.44$, $p = .08$). Although participants cooperated slightly more under time pressure than when given no instructions, this difference was not significant ($B = -.27$, $SE = 0.15$, $Z = -1.79$, $p = .17$).

To confirm the robustness of our time pressure findings, we conducted three additional analyses. First, modeling the effect of time pressure vs. time delay only within the outgroup conditions, we found a significant effect of time pressure ($Z = -3.60$, $p < .001$) such that participants were more likely to cooperate with outgroup members under time pressure (61%) than under time delay (43%). Second, modelling the effect of time pressure vs. time delay only within the losses frame conditions, we found a significant effect of time pressure ($Z = -3.68$, $p < .001$) such that participants

were more likely to cooperate in loss frames under time pressure (60%) than under time delay (41%). Finally, we modelled the effect of time pressure vs. time delay only within the outgroup and losses frame conditions. Again, we found a significant effect of time pressure ($Z = -3.63$, $p < .001$) such that participants were more likely to cooperate under time pressure (62%) than under time delay (35%), even towards outgroup members and even in a losses frame.

Overall, then, our analysis of simple effects demonstrates that, compared to time delay, time pressure leads to greater cooperation even in specific conditions where we hypothesized it might break down. Specifically, time pressure increases cooperation even for outgroups with a difference of 18 percentage points; just when analyzing losses frames with a difference of 19 percentage points; and especially for outgroup members under a losses frame, with a difference of 27 percentage points.

Discussion

In this study we investigated intuitive cooperation, considering both methodological and theoretical challenges to this work. Our findings make three main contributions to the literature. First, our work addresses concerns about selection bias – the asymmetric exclusion of high numbers of participants across conditions – that have plagued work on intuitive cooperation (Bouwmeester et al., 2017; Tinghög et al., 2013). Specifically, we drastically reduced participant exclusions from around 50% to just 2%, equated exclusion rates across the two active manipulation conditions, and ensured that all participants understood the structure of the game. Even when addressing these methodological issues we observed greater cooperation in the time pressure condition relative to the time delay condition, supporting the claims of the SSH. At the same time, however, by including a neutral control condition, we found that intuitive cooperation effects seem driven more by time-delay induced deliberation reducing cooperation than by time-pressure induced intuition increasing cooperation.

Second, our work sheds light on intuitive cooperation in intergroup

contexts. Contrary to traditional assumptions about implicit intergroup cognition, according to which people harbour intuitive biases that are particularly revealed in the absence of deliberative control, increased time pressure did not detectably increase or decrease intergroup bias. In other words, that is, we did not detect a two-way interaction between time pressure and group status on cooperation. We did employ an arbitrary, competitive groups paradigm in the present study and so it remains possible that a more forceful manipulation drawn from real social groups would yield a different result (as seen in Rand et al., 2015). Yet, note that our group manipulation was strong enough to yield a significant main effect of group membership: People cooperated more with members of their own group across all conditions. Critically, group membership also did not eliminate the effect of intuition versus deliberation on cooperation. Just as intuition enhances cooperation with ingroup members, we find that it also enhances cooperation with outgroup members. In sum, then, intuition increases prosociality but also fails to eradicate bias: Under time pressure participants cooperate with everybody more, but regardless of time pressure they cooperate with ingroups most of all. While we may be intuitive cooperators, we are not intuitively impartial.

Third, by investigating the boundary condition of losses framing on the effects of time pressure, ours is among the first studies to look at intuitive cooperation in both gains and losses frames. We extended previous work by showing that time pressure encourages cooperation even in losses frames, and even towards outgroup members in losses frames, and that there was greater overall cooperation in the gains context relative to the losses. This is consistent with work supporting a self-oriented perspective on loss aversion, whereby an aversion to losses makes people demand more in interpersonal interactions (De Dreu, Emans, Vliert, & Carnevale, 1994; De Dreu, Emans, & Van de Vliert, 1992), increases cheating (Schindler & Pfattheicher, 2017), and reduces prosocial behaviour (Brewer & Kramer, 1986). On the other hand, this is perplexing because it contrasts with more recent work

taking an other-oriented perspective on loss aversion (e.g. Everett, Faber, & Crockett, 2015; Leliveld, Beest, Dijk, & Tenbrunsel, 2009) in which a loss frame encourages greater prosocial behaviour. This may reflect the cognitive complexity of our loss-framed PD, where participants had to remember that they had a starting payment of \$0.80 and then mentally subtract the values in the matrix from that. If this complexity prompted greater deliberation, this could result in reduced cooperation simply because deliberative thinking promotes defection in the PD. This remains an interesting direction for further study.

At least two limitations of this work bear mention. First, we used deception in this task – both in the group manipulation and by not paying participants according to their decisions. This is one feature that differentiates our work from previous studies (e.g. Bouwmeester et al., 2017) and it would be preferable in future work for a full non-deception procedure to be implemented. Second, while we had sufficient power to moderately sized main effects, and also to detect the potential elimination of the “intuitive cooperation” effect in the critical cells of our design, we had weak statistical power to detect moderate two-interactions among variables, and insufficient statistical power to reliably detect moderate three-way interactions. Thus, our research provides strong evidence that intuitive cooperation effects exist even towards outgroup members, in a loss frame, and under the combination of these factors. It does not, however, provide strong evidence regarding whether the effects that exist are weaker or stronger than in the standard ingroup/gain frame condition. It would be interesting for future work to use a larger sample to explore such interaction effects.

In conclusion, we studied the effects of intuition versus deliberation on cooperation in an intergroup context, developing a new procedure that addresses methodological concerns that have plagued previous work. Our results show that both cooperation and in-group favoritism emerge even when not given the chance to deliberate, suggesting that while we might be intuitive cooperators, we are not intuitively impartial.

References

- Artavia-Mora, L., Bedi, A. S., & Rieger, M. (2017). Intuitive help and punishment in the field. *European Economic Review*, 92, 133–145. <https://doi.org/10.1016/j.euroecorev.2016.12.007>
- Bear, A., Kagan, A., & Rand, D. G. (2017). Co-evolution of cooperation and cognition: the impact of imperfect deliberation and context-sensitive intuition. *Proc. R. Soc. B*, 284(1851), 20162326. <https://doi.org/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(4), 936–941. <https://doi.org/10.1073/pnas.1517780113>
- Bouwmeester, S., Verkoeijen, P. P. J. L., Aczel, B., Barbosa, F., Bègue, L., Brañas-Garza, P., ... Wollbrant, C. E. (2017). Registered Replication Report: Rand, Greene, and Nowak (2012). *Perspectives on Psychological Science*. <https://doi.org/10.1177/1745691617693624>
- Brewer, M. B., & Kramer, R. M. (1986). Choice behavior in social dilemmas: Effects of social identity, group size, and decision framing. *Journal of Personality and Social Psychology*, 50(3), 543–549. <https://doi.org/10.1037/0022-3514.50.3.543>
- Capraro, V., Jordan, J. J., & Rand, D. G. (2014). Heuristics guide the implementation of social preferences in one-shot Prisoner's Dilemma experiments. *Scientific Reports*, 4, 6790. <https://doi.org/10.1038/srep06790>
- Chang, L. W., Krosch, A. R., & Cikara, M. (2016). Effects of intergroup threat on mind, brain, and behavior. *Current Opinion in Psychology*, 11, 69–73. <https://doi.org/10.1016/j.copsy.2016.06.004>
- Cikara, M., Jenkins, A. C., Dufour, N., & Saxe, R. (2014). Reduced self-referential neural response during intergroup competition predicts competitor harm. *NeuroImage*, 96(1), 36–43. <https://doi.org/10.1016/j.neuroimage.2014.03.080>
- Cikara, M., & Van Bavel, J. J. (2014). The neuroscience of intergroup relations: An integrative review. *Perspectives on Psychological Science*, 9(3), 245–274. <https://doi.org/10.1177/1745691614527464>
- Cone, J., & Rand, D. G. (2014). Time pressure increases cooperation in competitively framed social dilemmas. *PLOS ONE*, 9(12), e115756. <https://doi.org/10.1371/journal.pone.0115756>
- Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2002). The police officer's dilemma: using ethnicity to disambiguate potentially threatening individuals. *Journal of Personality and Social Psychology*, 83(6), 1314. <https://doi.org/10.1037/0022-3514.83.6.1314>
- De Dreu, C. K. W., Dussel, D. B., & Velden, F. S. T. (2015). In intergroup conflict, self-sacrifice is stronger among pro-social individuals, and parochial altruism emerges especially among cognitively taxed individuals. *Frontiers in Psychology*, 6. <https://doi.org/10.3389/fpsyg.2015.00572>
- De Dreu, C. K. W., Emans, B. J. M., Vliert, E., & Carnevale, P. J. (1994). Effects of gain-loss frames in negotiation: Loss aversion, mismatching, and frame adoption. *Organizational Behavior and Human Decision Processes*, 60, 90–107. <https://doi.org/10.1006/obhd.1994.1076>
- De Dreu, C. K. W., Emans, B. J., & Van de Vliert, E. (1992). Frames of reference and cooperative social decision-making. *European Journal of Social Psychology*, 22(3), 297–302. <https://doi.org/10.1002/ejsp.2420220307>
- De Dreu, C. K. W., & McCusker, C. (1997). Gain-loss frames and cooperation in two-person social dilemmas: A transformational analysis. *Journal of Personality and Social Psychology*, 72(5), 1093. <https://doi.org/10.1037/0022-3514.72.5.1093>

- de Heus, P., Hoogervorst, N., & Dijk, E. van. (2010). Framing prisoners and chickens: Valence effects in the prisoner's dilemma and the chicken game. *Journal of Experimental Social Psychology*, 46(5), 736–742. <https://doi.org/10.1016/j.jesp.2010.04.013>
- Dunham, Y., Baron, A. S., & Banaji, M. R. (2008). The development of implicit intergroup cognition. *Trends in Cognitive Sciences*, 12(7), 248–253. <https://doi.org/10.1016/j.tics.2008.04.006>
- Everett, J. A. C., Faber, N. S., & Crockett, M. J. (2015). The influence of social preferences and reputational concerns on intergroup prosocial behaviour in gains and losses contexts. *Royal Society Open Science*, 2(12), 150546. <https://doi.org/10.1098/rsos.150546>
- 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>
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: the implicit association test. *Journal of Personality and Social Psychology*, 74(6), 1464–1480. <https://doi.org/10.1037/0022-3514.74.6.1464>
- 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. *Psychological Science*, 19(4), 405–411. <https://doi.org/10.1111/j.1467-9280.2008.02100.x>
- Hewstone, M., Rubin, M., & Willis, H. (2002). Intergroup bias. *Annual Review of Psychology*, 53(1), 575–604. <https://doi.org/10.1146/annurev.psych.53.100901.135109>
- Komorita, S. S., & Carnevale, P. (1992). Motivational arousal vs. decision framing in social dilemmas. In W. B. G. Liebrand, D. M. Messick, & H. A. M. Wilke (Eds.), *Social dilemmas: Theoretical issues and research findings* (pp. 209–223). Elmsford, NY, US: Pergamon Press.
- Lai, C. K., Skinner, A. L., Cooley, E., Murrar, S., Brauer, M., Devos, T., ... Nosek, B. A. (2016). Reducing implicit racial preferences: II. Intervention effectiveness across time. *Journal of Experimental Psychology: General*, 145(8), 1001–1016. <https://doi.org/10.1037/xge0000179>
- Leliveld, M. C., Beest, I. van, Dijk, E. van, & Tenbrunsel, A. E. (2009). Understanding the influence of outcome valence in bargaining: A study on fairness accessibility, norms, and behavior. *Journal of Experimental Social Psychology*, 45(3), 505–514. <https://doi.org/10.1016/j.jesp.2009.02.006>
- Ma, Y., Liu, Y., Rand, D. G., Heatherton, T. F., & Han, S. (2015). Opposing oxytocin effects on intergroup cooperative behavior in intuitive and reflective minds. *Neuropsychopharmacology*, 40(10), 2379–2387. <https://doi.org/10.1038/npp.2015.87>
- Miller, S. L., Zielaskowski, K., & Plant, E. A. (2012). The basis of shooter biases: Beyond cultural stereotypes. *Personality and Social Psychology Bulletin*, 38(10), 1358–1366. <https://doi.org/10.1177/0146167212450516>
- Rand, D. G. (2016). Cooperation, fast and slow: Meta-analytic evidence for a theory of social heuristics and self-interested deliberation. *Psychological Science*, 27(9), 1192–1206. <https://doi.org/10.1177/0956797616654455>
- Rand, D. G., Brescoll, V. L., Everett, J. A. C., Capraro, V., & Barcelo, H. (2016). Social heuristics and social roles: Intuition favors altruism for women but not for men. *Journal of Experimental Psychology: General*, 145(4), 389–396. <https://doi.org/10.1037/xge0000154>
- Rand, D. G., & Epstein, Z. G. (2014). Risking Your Life without a Second Thought: Intuitive Decision-Making and Extreme Altruism. *PLOS ONE*, 9(10), e109687. <https://doi.org/10.1371/journal.pone.0109687>
- Rand, D. G., Greene, J. D., & Nowak, M. A. (2012). Spontaneous giving and calculated greed. *Nature*, 489(7416),

- 427–430.
<https://doi.org/10.1038/nature11467>
- Rand, D. G., Greene, J. D., & Nowak, M. A. (2013). Rand et al. reply. *Nature*, 498(7452), E2–E3.
<https://doi.org/10.1038/nature12195>
- Rand, D. G., Newman, G. E., & Wurzbacher, O. M. (2015). Social context and the dynamics of cooperative choice. *Journal of Behavioral Decision Making*, 28(2), 159–166.
<https://doi.org/10.1002/bdm.1837>
- Rand, D. G., Peysakhovich, A., Kraft-Todd, G. T., Newman, G. E., Wurzbacher, O., Nowak, M. A., & Greene, J. D. (2014). Social heuristics shape intuitive cooperation. *Nature Communications*, 5, 3677.
<https://doi.org/10.1038/ncomms4677>
- Schindler, S., & Pfattheicher, S. (2017). The frame of the game: Loss-framing increases dishonest behavior. *Journal of Experimental Social Psychology*, 69, 172–177.
<https://doi.org/10.1016/j.jesp.2016.09.009>
- Sherif, M. (1966). *In common predicament: Social psychology of intergroup conflict and cooperation*. Boston, MA: Houghton Mifflin.
- Ten Velden, F. S., Daughters, K., & De Dreu, C. K. W. (2017). Oxytocin promotes intuitive rather than deliberated cooperation with the in-group. *Hormones and Behavior*.
<https://doi.org/10.1016/j.yhbeh.2016.06.005>
- Tinghög, G., Andersson, D., Bonn, C., Böttiger, H., Josephson, C., Lundgren, G., ... Johannesson, M. (2013). Intuition and cooperation reconsidered. *Nature*, 498(7452), E1–E2.
<https://doi.org/10.1038/nature12194>
- Verkoeijen, P. P. J. L., & Bouwmeester, S. (2014). Does intuition cause cooperation? *PLOS ONE*, 9(5), e96654.
<https://doi.org/10.1371/journal.pone.0096654>
- Wright, P. (1974). The harassed decision maker: Time pressures, distractions, and the use of evidence. *Journal of Applied Psychology*, 59(5), 555–561.
<https://doi.org/10.1037/h0037186>
- Zaki, J., & Mitchell, J. P. (2013). Intuitive prosociality. *Current Directions in Psychological Science*, 22(6), 466–470.
<https://doi.org/10.1177/0963721413492764>

Tables

Table 1. Rates of cooperation across conditions.

Gains Frame						
	In-group			Out-group		
	Cell N	Freq. of Cooperation	Percentage Cooperating	Cell N	Freq. of Cooperation	Percentage Cooperating
Intuition	90	62	69%	99	60	61%
Reflection	96	60	62%	96	48	50%
Control	117	79	68%	102	57	56%

Losses Frame						
	In-group			Out-group		
	Cell N	Freq. of Cooperation	Percentage Cooperating	Cell N	Freq. of Cooperation	Percentage Cooperating
Intuition	84	49	58%	93	58	62%
Reflection	88	41	47%	86	30	35%
Control	98	51	52%	91	45	49%

Figures



Figure 1. Example Matrices given to participants (left: gains frame with in-group partner; right: losses frame with out-group partner)

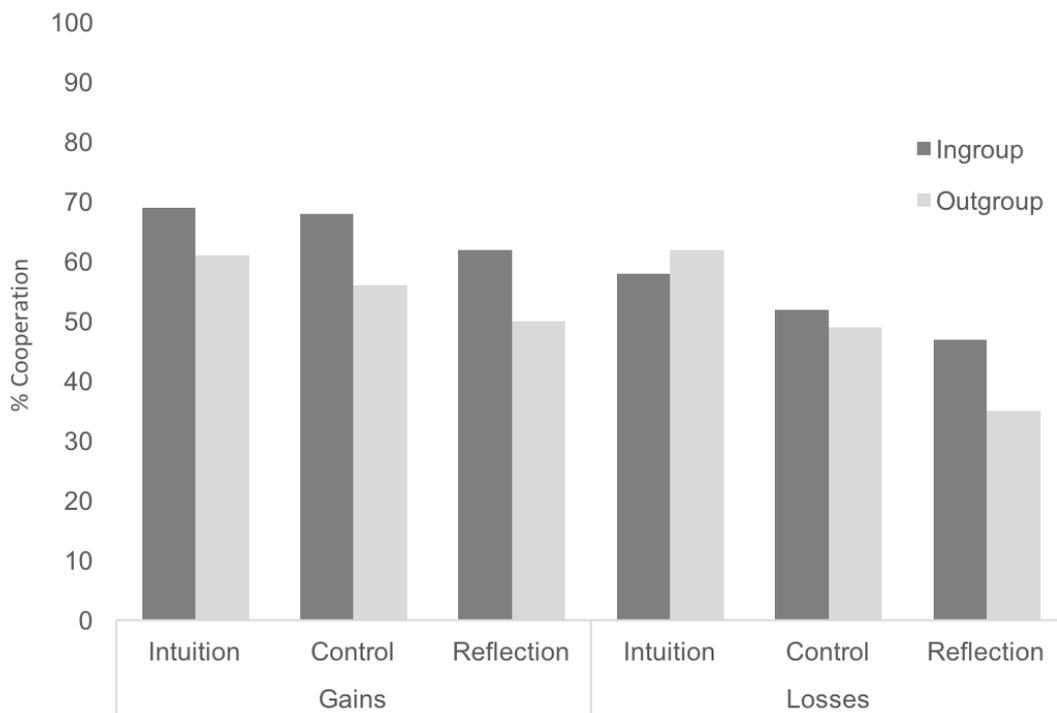


Figure 2. Rates of cooperation across conditions.

¹ While 1687 participants opened the study link, only 1626 moved beyond the first page of the personality questions for the team assignment (dropout of 4%) and only 1316 moved beyond the instructions of the game and comprehension checks (dropout of 18%). This led to a total dropout rate of 22%, but one that was broadly equivalent across conditions. As can be seen in Table 1, there were roughly the same number of people in both the in-group ($N = 573$) and out-group ($N = 567$) conditions, and both the time delay ($N = 366$) and time pressure ($N = 366$) conditions. There were slightly more participants in the control condition ($N = 407$) compared to the time delay and time pressure, and slightly more participants in the gains conditions ($N = 600$) than the losses conditions ($N = 540$), but these differences were small.