



If not me then we: Goal tradeoffs in decision-making for the self, ingroup, and outgroup[☆]

Suraiya Allidina^{a,*}, William A. Cunningham^b

^a Ohio State University, 1827 Neil Ave, Columbus, OH 43210, USA

^b University of Toronto, 100 St George St, Toronto, ON M5S 3G3, Canada

ARTICLE INFO

Keywords:

Intergroup behaviour
Goal trade-offs
Multiple goal pursuit
Prosociality

ABSTRACT

Navigating the social world requires individuals to balance multiple goals, including the drives to improve one's own outcomes, aid ingroup members, and help or hurt outgroup members. While self-interest and intergroup bias are both well-established motivational phenomena, less is known about how these goals may interact. Here we examine the nature of goal tradeoffs in intergroup decision-making using a novel task in which participants simultaneously make monetary decisions for themselves, an arbitrary ingroup, and the corresponding outgroup. Across four behavioural studies and one eye-tracking study (total $N = 704$), we find that goals in intergroup contexts are pursued sequentially rather than concurrently, with non-linear upweighting of group-related goals when self-related goals cannot be pursued. Further, we find evidence for stronger self-ingroup than self-outgroup tradeoffs, which manifest in both altered attention to information and altered use of the attended information in decision-making. The results shed light on the cognitive structuring of interrelated goals in intergroup decision-making, furthering our understanding of when and how both intergroup biases and prosocial behaviour may emerge.

Social decision-making involves a collection of interconnecting goals and values: people are highly attuned to groups, generally opting to favour their ingroup members when possible, but are also motivated to prioritize themselves. While tendencies towards self-interest and intergroup bias are well-established, these motivations do not exist in isolation. In particular, intergroup contexts with no self-related implications are rare, as are situations involving the self and nameless others with no group memberships. Rather, there are often multiple relevant social recipients, each with their own group identities and memberships, and thus multiple proximal goals for the perceiver to balance. Instead of simply proceeding in parallel, these various interconnected goals are likely to influence one another, giving rise to tradeoffs in goal pursuit such that people must dynamically choose which goals to prioritize. In the current work, we examine the nature of these tradeoffs, asking how social decision-makers prioritize different self- and group-related goals in intergroup contexts.

1. Motivations for self- and group-interest

Humans are fundamentally group-based: across contexts,

individuals, and group types, the tendency to care more about the outcomes of one's own group members predominates. When given the chance, people will help ingroup members over outgroup members (Tajfel, Billig, & Bundy, 1971; de Dreu, Dussel, & Ten Velden, 2015; Balliet, Wu, & Dreu, 2014; Fehr, Bernhard, & Rockenbach, 2008), and will sometimes even actively prefer to hurt members of other groups simply because of their group membership (e.g., displaying pleasure in response to the pain of outgroup members; Leach, Spears, Branscombe, & Doosje, 2003; Smith, Powell, Combs, & Schurtz, 2009; Cikara, Botvinick, & Fiske, 2011). Preferences for ingroup members persist even when groups are completely arbitrary and based on relatively meaningless distinctions (Ahmed, 2007; Tajfel et al., 1971; Halevy, Weisel, & Bornstein, 2011; Kramer & Brewer, 1984), with motivations towards these minimal groups paralleling the motivations people have towards groups divided by more meaningful dimensions. In parallel with these group-oriented motives, people are also motivated to maintain their own positive outcomes. In fact, some accounts of intergroup bias explain ingroup-oriented behaviour as an outcome of self-oriented motivations, with the ingroup viewed as an extension of the self (Purzycki & Lang, 2019; Swann, Gómez, Seyle, Morales, & Huici, 2009). Regardless of the

[☆] This paper has been recommended for acceptance by Paul Conway.

* Corresponding author.

E-mail address: allidina.1@osu.edu (S. Allidina).

ultimate causes of ingroup-oriented behaviour, however, in the proximal sense people must balance their goals to benefit themselves and their goals to benefit their groups (often at the expense of other groups).

Drives towards self-interest and group-interest do not exist in isolation: helping the ingroup may sometimes be beneficial to the self but at other times may require self-sacrifice. With both self- and group-related goals featuring prominently in social decision-making, tradeoffs may often need to be made and some goals pursued at the expense of others. For example, individuals who choose to join a demonstration to protest the unfair treatment of their group members must weigh the benefit to their group against the personal risk to themselves that attending the protest may bring, potentially choosing to sacrifice their goal of self-preservation to better benefit their group. Similarly, one study found that Black college students almost universally chose to donate money to Obama rather than Romney in the 2012 presidential election (pursuing the goal aligned with their group), but this was significantly reduced when they could personally earn money by donating to Romney (with about 30% choosing to prioritize self-interest over group interests; White, Laird, & Allen, 2014). As these examples demonstrate, progress towards one goal may depend on opportunities to advance other goals. Initial research supports this idea, suggesting that ingroup-oriented prosocial behaviours may be reduced when self-interest features more prominently (Stagnaro, Dunham, & Rand, 2018). A larger body of work has explored how people's chosen priorities relate to one another, asking whether those who are more prosocial (i.e., who favour the "other" in self-other decisions) are more or less likely to display parochial altruism (i.e., to favour the ingroup in ingroup-outgroup decisions) (Aaldering, Ten Velden, van Kleef, & De Dreu, 2018; Abbink, Brandts, Herrmann, & Orzen, 2012; Böhm, Rusch, & Baron, 2018; de Dreu, 2010; de Dreu et al., 2015; Thielmann & Böhm, 2016; Yamagishi & Mifune, 2016). However, while this research explores the co-occurrence of these motives, their interaction has been relatively underexplored, and it remains an open question how behaviour towards one goal (e.g., helping the ingroup) changes depending on another goal (e.g., helping the self).

2. Concurrent vs. sequential goal pursuit

Given these various often conflicting drives, how might ingroup- and outgroup-oriented behaviour change as a function of self-interest? To answer this question, we turn to theories of multiple goal pursuit, which suggest that perceivers who wish to make progress towards several goals can do so either *concurrently* or *sequentially* (Orehek & Vazeou-Nieuwenhuis, 2013; Kruglanski et al., 2013). Concurrent goal pursuit involves simultaneously making progress towards multiple active goals at once, looking for actions that fulfill multiple goals rather than just one. Since multiple goals are being attained simultaneously, each individual goal will have slightly slower goal progress than if it were the only goal being pursued. Thus, concurrent goal pursuit is likely when multiple goals have approximately equal importance for the goal pursuer. Sequential goal pursuit, on the other hand, involves more heavily prioritizing those goals which are seen as more important. Resources are primarily allocated towards one goal at a time, with the actor unwilling to make even small sacrifices to the primary goal to progress towards secondary goals. This approach to multiple goal pursuit is more likely when a single goal is seen as much more important than other goals, with the primary goal "shielding" the other goals from goal-relevant actions (Shah, Friedman, & Kruglanski, 2002). Critically, these two possibilities represent two ends on a spectrum rather than two dichotomous strategies.

To understand how self- and group-related goals are prioritized in intergroup decision-making contexts, we must examine how progression towards group-related goals changes as a function of self-interest. Thus, we need to examine the pursuit of group-related goals when self-related goals are vs. are not also active. Concurrent and sequential approaches to goal pursuit would produce different effects on goal-directed behaviour under these conditions. When multiple goals are active and pursued

concurrently, the agent works towards all of them simultaneously. Goals that are more important are given higher weight than goals that are less important, but all goals are weighted in a linear manner and progressed towards simultaneously. Here, the difference between pursuing a single goal on its own and pursuing the same goal alongside other goals should not be too drastic, since the individual is still working towards the same goal in both cases. In contrast, sequential goal pursuit would produce a more drastic difference between pursuing a single goal on its own and pursuing it while other goals are also active. When the primary goal is active, behaviour should be mainly focused on attaining it; secondary goals should therefore see a larger upweighting when this primary goal is not active. Here, goals are prioritized in a non-linear manner such that the mere presence or absence of opportunities to progress towards a primary goal determines progression towards secondary goals.

Thus, in a context in which participants are primarily motivated by self-interest but also motivated to help their ingroups and potentially help or hurt the outgroups, these two possibilities make differing predictions. If goals for the self and for groups are fulfilled sequentially, we should see a more drastic increase in progression towards group-related goals when the self-related goals are blocked. In other words, people should primarily help themselves when it is possible to do so, and only start considering others more when it is not possible to help themselves. In this case, the more important variable should be whether or not the self-related goal can be pursued. If, on the other hand, these goals are pursued concurrently, group-oriented goals would still be upweighted when self-related goals are blocked, but only in a linear manner; there should be nothing special about the absence of opportunities to pursue self-related goals. Rather, the value of the potential outcomes for the self should linearly influence participants' behaviour towards the ingroup and outgroup. We highlight once again that these possibilities represent ends on a spectrum, and processing that is wholly sequential (e.g., with complete inattention to the groups until the self is irrelevant) is unlikely. Thus, we examine whether evidence for a significant degree of sequential processing is present.

The presence of sequential as opposed to continuous goal pursuit may be determined by the relative primacy of one goal over others. In particular, identities at the level of the individual self (rather than the self as a group member, or "collective self") seem to be motivationally primary. Research by Gaertner and colleagues (Gaertner et al., 2012; Gaertner, Sedikides, & Graetz, 1999; Gaertner, Sedikides, Vevea, & Iuzzini, 2002) demonstrates that motivations for the individual self tend to dominate: participants consider a threat to their individual selves more severe than one to their collective selves, with more negative reactions and compensatory mechanisms elicited by the former. This is the case even when motivations for the ingroup are quite strong and participants' own identities are fused with those of their ingroups (Heger, Voorhees, Porter, & Gaertner, 2023). It is possible that such motivational primacy will result in sequential pursuit of goals for the self and for group members, with positive outcomes for the self pursued when possible and group outcomes pursued when not. Thus, by examining the simultaneous pursuit of self- and group-related goals here, we aim to reveal some of the mechanisms through which the motivational prioritization of social information occurs.

3. Previous research

Previous research has thoroughly examined motives for self-interest and for intergroup bias in isolation, but existing paradigms are less well-suited to examining the precise nature and dynamics of the tradeoffs between these goals. Existing economic games that have been used to examine self- and group-related monetary goals fall into three broad classes. First are social dilemma games, in which a participant is given an allocation of money which they can choose to either keep for themselves or contribute to various "pools" that serve different functions like benefitting the ingroup, benefitting the ingroup while also hurting the outgroup, or benefitting everyone regardless of group membership

(Halevy, Bornstein, & Sagiv, 2008; Aaldering et al., 2018; Aaldering & Böhm, 2020). Second are intergroup competition games, where participants can choose whether or not to invest a sum of money to their ingroup in order to beat the other group (e.g., Abbink et al., 2012). Last are single-sided games where the participant can unilaterally decide how much of a pool to give to another person, varying in group membership, from their own monetary allocation (Purzycki & Lang, 2019; Rahal, Fiedler, & De Dreu, 2020).

Overall, these paradigms have demonstrated motivations to help the ingroup more than the outgroup and, in some cases, to also hurt the outgroup. However, they are not as well-suited to exploring the question posed here, as they do not allow us to compare group-oriented motives under both the presence and absence of self-oriented motives (the critical comparison for this research question). Some paradigms, like the intergroup competition games, deliberately conflate self- and ingroup-interest materially, such that contributing to the ingroup-related pool also helps the self (e.g., Abbink, Brandts, Herrmann, & Orzen, 2010, Abbink et al., 2012). These tasks are therefore not suitable for examining tradeoffs between these goals, as self-interested and ingroup-directed behaviour cannot be easily separated. Further, all the paradigms typically consist of one or a very small number of decisions and payoff structures for a given participant. Thus, comparison of goal-directed behaviour under different goal contingencies is not possible, as each participant often sees only a single set of outcomes. For example, the intergroup dictator game involves directly incurring a cost to the self to help another (Stagnaro et al., 2018), which directly confounds self- and other-related goals. A decision to prioritize the self in these situations could be driven by a high degree of self-interest, a low degree of other-interest, or some combination thereof. Since self- and other-relevant outcomes are not varied separately, it is not possible to separate out the effects of self- and other-regarding preferences.

In contrast, for this research we need a paradigm that allows participants to actively work towards goals they may have for the self, the ingroup, and the outgroup in a relatively orthogonal manner. For some decisions these goals should be aligned, while for others they should be in conflict. In this way, we can examine how behaviour changes under various goal contingencies. For example, the paradigm should allow us to compare how much participants help the self when doing so hurts the ingroup compared to when doing so helps the ingroup.

4. The current research

Gaining a fuller picture of social and intergroup motivations requires us to examine these goals in tandem. In the current work, we therefore investigate goal tradeoffs by examining the attentional dynamics of information selection and integration over time. Across four behavioural studies and one eye-tracking study, we use a novel paradigm to examine how people resolve tradeoffs between goals for themselves, an arbitrary

ingroup, and the corresponding outgroup. Participants are first randomly assigned to one of two minimal groups before completing a decision-making task in which they choose monetary outcomes for themselves, for other members of their ingroup, and for members of their outgroup. On each trial of this task, participants are presented with two potential point allocations that they can choose between. Each option contains a monetary outcome for themselves and a monetary outcome for a group, with one option containing an ingroup outcome and the other containing an outgroup outcome (see Fig. 1). For example, participants may choose between getting 10 points for themselves along with 50 points for their ingroup members, or 5 points for themselves along with 50 points for their outgroup members.

This paradigm overcomes the limitations of other decision-making tasks outlined above, with a number of key features that allow us to answer our research questions. First, goals for the self, ingroup, and outgroup are fully separable and relatively orthogonal in this task. As each potential outcome seen on screen varies independently on each trial, multiple types of goal contingencies are present on different trials. Critically, self and ingroup interests are not materially conflated in this task; outcomes participants choose for the ingroup affect all other members of the participant's group but not their own outcomes. Some trials involve relatively easy choices in which self-interest and ingroup-interest are aligned (e.g., the favourable option A contains +10 points for the self and +50 points for the ingroup, whereas the unfavourable option B contains +1 point for the self and +50 points for the outgroup). Other trials involve tradeoffs between goals (e.g., for a participant who is motivated to help themselves and their ingroup members, option A contains +10 points for the self and -50 points for the ingroup, whereas option B contains +1 point for the self and -50 points for the outgroup, forcing participants to choose whether to help themselves at a cost to their ingroup members). Finally, other trials are irrelevant to self-related goals and only group-related goals are active (e.g., option A contains +5 points for the self and +50 points for the ingroup, whereas option B contains +5 points for the self and +50 points for the outgroup). Critically, information about the self, ingroup, and outgroup is present on each trial in the same choice structure, as we aim to examine how people navigate tradeoffs between all three of these goals (though we separate ingroup- and outgroup-related choices to ensure the robustness of these effects in study 2).

Second, participants in this paradigm can materially affect each other's outcomes. Participants are informed that their final point totals are influenced by 1) the points they get for themselves in the task, 2) the points that the other members of their ingroup got for the group, and 3) the points that members of the outgroup got for them. For example, a given participant on the red team has a final point total determined by a) the points the participant got for themselves, b) the points the other members of the red team got for the ingroup, and c) the points the members of the blue team got for the outgroup. In this way, participants




Fig. 1. Example trial for Studies 1a-1c (A) and Study 2 (B). In all studies, participants had to choose to take either both options on the left or both options on the right. Each option contained an outcome for themselves and an outcome for either their ingroup or outgroup. In studies 1a-1c, one option contained an option for the ingroup and the other contained an option for the outgroup, whereas outcomes for the ingroup and outgroup were separated across blocks in study 2.

are interdependent: their choices affect others, and others' choices affect them. Critically, points for the self and for the ingroup are fully separated in this design: the points that a given participant earns for their ingroup members do not directly contribute to the participant's own monetary outcomes. Given the large body of research demonstrating highly similar results when monetary choices are hypothetical or real (Ben-Ner, Kramer, & Levy, 2008; Camerer, 1989; Gillis & Hettler, 2007; Kühberger, Schulte-Mecklenbeck, & Perner, 2002), studies 1a-1c used purely hypothetical incentives. However, to confirm that the hypothetical incentives did not alter participants' behaviour, studies 2 and 3 used real monetary incentives, in which each person could earn a bonus based on both their own behaviour and other participants' choices (as described above). Results replicate in these studies, indicating that participants' behaviour did not significantly change when incentives were hypothetical or real.

Finally, no meaningful distinction exists between the two groups in this task and there is no monetary incentive to help the ingroup or hurt the outgroup. Participants can gain or lose points for their ingroup and outgroup members, but the total points that they earn for others has no bearing on their own outcomes, especially since other participants are not aware of their choices and no opportunity for direct reciprocity is therefore present (although expectations of reciprocity may still drive decisions; Allidina, Arbuckle, & Cunningham, 2019; Gaertner & Insko, 2000; Rabbie, Schot, & Visser, 1989; Rabbie & Lodewijkx, 1994). However, given the decades of research showing the powerful effects of completely arbitrary minimal groups on attitudes, beliefs, and behaviour, we expect that participants will nonetheless choose to favour the ingroup and potentially even actively hurt the outgroup. The use of minimal groups in this context is key for isolating the motivational effects of interest, allowing us to more fully dissociate effects of self- and group-related goals on behaviour. Since motivations around real-world social groups can be more complex and multifaceted (including moral concerns that may exist around the use of a person's group membership when allocating money), minimal groups provide an alternative environment which contains the powerful group-based motivations of interest but minimizes other secondary or distracting forces.

Here, we use this paradigm to explore goal prioritization in intergroup contexts, aiming to arbitrate between the two possibilities outlined above for navigating goal conflict. On one hand, people may weight information through a simple linear combination in which they consider all relevant information and combine it according to their priorities to make a decision (indicative of concurrent goal pursuit). On the other hand, they may prioritize goals in a non-linear manner, such that the mere presence or absence of opportunities to progress towards a primary goal is more relevant (indicative of sequential goal pursuit). We arbitrate between these possibilities by examining decisions to help or hurt the ingroup and outgroup when the self goal is active (i.e., the choice matters for the self, because one side has a better outcome for the self than the other side) compared to when the self goal is inactive (i.e., the choice does not matter for the self, because both self outcomes are equal). If people are progressing towards multiple goals concurrently and thus weighting information linearly, the value of the potential outcomes for the self should linearly influence behaviour towards the ingroup and the outgroup. If, however, people are progressing towards goals sequentially and thus non-linearly prioritizing certain goals over others, we should see distinct differences between the use of group-related information when self outcomes are equal and unequal, with differences between other values of self outcome differences less pronounced.

Across 5 studies, we therefore investigate how the ability to directly fulfill self-related goals might shape the fulfillment of group-related goals by comparing the use of group-related information at various levels of self outcome differences. Studies 1a-1c examine this phenomenon using the task described above, in which people simultaneously make decisions for the self, ingroup, and outgroup. To rule out alternative explanations, Study 2 replicates our findings using a variation of

the task in which decisions for the ingroup and outgroup are separated across blocks and monetary incentives are real. Finally, study 3 turns to the mechanisms of this goal prioritization, using eye-tracking to examine how information is selectively attended to and prioritized for decision-making. By investigating the nature and degree of social goal tradeoffs in this manner, we aim to further our understanding of when and how both intergroup biases and prosocial behaviour will emerge. This research was approved by the university's Research Ethics Board, and all relevant ethical regulations were complied with. All studies, manipulations, measures, and data exclusions are reported, and sample sizes were determined before any data analysis for each study.

5. Studies 1a-1c

We first ran a series of three behavioural studies to begin examining the tradeoffs people make when balancing multiple social goals. As these three studies had highly similar methodologies, serving largely as replications, we report them together for easier comparison of results across studies. The critical question for this set of studies concerns how people navigate tradeoffs between their proximal goals for themselves and for others, including potential goals to help the ingroup and to hurt the outgroup. We examine whether people pursue self- and group-related goals *concurrently*, such that information is weighted linearly and combined for decision-making, or *sequentially*, such that the mere presence or absence of opportunities to pursue the primary self-related goal shapes pursuit of group-related goals.

5.1. Methods

5.1.1. Participants

Study 1a. 108 undergraduate students (45 East/Southeast Asian, 35 White, 16 South Asian, 1 Hispanic, 1 Native American, 10 other races/ethnicities; 60 male, 48 female; mean age = 19.4, SD age = 2.3) were recruited from the University of Toronto and participated for partial course credit. 5 participants were excluded due to an error in stimulus presentation, leaving 103 participants for analysis.

Study 1b. 188 undergraduate students (118 East/Southeast Asian, 34 White, 18 South Asian, 4 Black, 1 Hispanic, 9 other races/ethnicities, 4 did not report race/ethnicity; 129 female, 55 male, 4 did not report gender; mean age = 20.2, SD age = 2.2) were recruited from the University of Toronto and participated in the study in exchange for partial course credit.

Study 1c. 300 participants (211 White, 23 Black, 19 East/Southeast Asian, 19 Hispanic, 5 South Asian, 4 Native American, 4 other races/ethnicities, 15 did not report race/ethnicity; 152 male, 133 female, 15 did not report gender; mean age = 36.1, SD age = 10.4) were recruited from Amazon's Mechanical Turk and paid for completing the study.

5.1.2. Procedure

In all three studies, participants were first randomly assigned to one of two minimal groups: the red team or the blue team. No information was given about these groups; they were simply told that there were two groups in the task and they were a member of the given team. They were informed that they would be playing a game in which they made decisions for themselves and for other people. Specifically, they would have to make various choices that would determine the points they got for themselves, the points they got for members of their assigned team (which did not directly affect their own points), and the points they got for members of the other team. They were then presented with six different cartoon avatars which varied in gender and racial appearance and were asked to choose the one that most resembled themselves to represent them in the game.

Participants then completed a monetary intergroup decision-making task. On each trial, they were presented with two options and would have to choose to take one or the other. Each option contained a point outcome for themselves. The self outcome on one side of the screen was

paired with an outcome for the ingroup, while the self outcome on the other side of the screen was paired with an outcome for the outgroup. Participants would have to choose to take either both options on the left or both options on the right. Possible values for the self outcomes were +1, +5, and +10. Possible values for the group outcomes were −50, −10, +10, +50; since participants were informed that each group consisted of 10 people, this translated into outcomes of −5 to +5 for each individual in the group (see Fig. 1 for an example trial). The task consisted of 80 trials in total. We expected that trials where participants had to choose between giving their ingroup a negative outcome and giving the outgroup a positive outcome would be especially difficult and revealing of participants' motivations. We therefore upsampled these types of trials, such that 40% of trials consisted of a negative ingroup outcome and a positive outgroup outcome. The remaining trials were split evenly between trial types, such that 20% of trials had a positive ingroup outcome and negative outgroup outcome, 20% had positive ingroup and outgroup outcomes, and 20% had negative ingroup and outgroup outcomes. Self outcomes were randomly determined on each trial, such that on average 33% of trials had two equal self outcomes.

By varying outcomes for the self and the groups trial by trial, this design allows for multiple types of decision trials with various degrees of goal conflict or alignment. On some trials, for example, a high outcome for the self would be accompanied by a negative outcome for the ingroup, such that in order for the participant to help themselves they must also harm the ingroup. A participant who is largely concerned with improving their own points and cares little about the ingroup may choose this outcome without much deliberation, while another who cares more about their group may choose to forego the benefit to themselves to avoid harming their group. On other trials, the two potential outcomes for the self might be equal, leaving only the ingroup and the outgroup information to base one's decision on. On these trials, a subject who is more motivated to help their ingroup may choose outcomes that are good for their team and avoid outcomes that are bad for their team, ignoring the other team's information altogether. In contrast, a subject who is specifically motivated to harm the other team may deliberately choose outcomes that are negative for the outgroup. Varying each kind of information across trials therefore allows us to investigate decisions under a variety of conflicting and/or aligning goals to examine how people trade off different goals.

The procedure of studies 1b and 1c were very similar to that of study 1a, with one addition. In the latter two studies, before completing the main block of trials described above, participants first completed a smaller block of 40 trials where they made simpler decisions between only two outcomes. Specifically, they would see two options on screen and had to choose one or the other. In study 1b, the options presented on either side of the screen varied between subjects as follows: self vs. self, ingroup vs. ingroup, self vs. ingroup, and self vs. outgroup. Study 1c had each of these combinations, as well as two additional: outgroup vs. outgroup and ingroup vs. outgroup. These blocks were included to allow participants to get accustomed to the task and to gain exploratory data for a different research question; no differences were found between conditions on behaviour in the subsequent main block, and these initial blocks are therefore not discussed further. After completing this initial block, all participants completed the same task as in study 1a.

Throughout these studies, we sought to maximize power through utilization of within-person designs analyzed using linear hierarchical models. To assess power, we conducted a sensitivity analysis to determine the smallest interaction estimate we could reliably detect using a z-test given our sample sizes, focusing on the interaction between the group outcome and whether or not the self outcomes were equal. This analysis indicated that after exclusions (described below), we had 80% power to detect interaction slope estimates as small as −0.14 in study 1a, −0.1 in study 1b, and −0.09 in study 1c.

5.1.3. Questionnaires

After completing the task, participants filled out information about

their demographics, and then were debriefed. Participants in studies 1b and 1c also completed the Social Dominance Orientation (SDO) scale, a 16-item scale measuring preferences for hierarchy and inequality among social groups (Pratto, Sidanius, Stallworth, & Malle, 1994). Participants responded to each item using a scale from 1 (strongly disagree) to 7 (strongly agree), and responses to all items were then averaged for analysis. See Supplementary Materials for analyses of questionnaires.

5.2. Results

5.2.1. Exclusions

We applied a number of exclusion criteria to filter out subjects who may not have been paying attention during the task. In particular, we removed subjects if >15% of their trials had a latency below 150 ms (suggesting they were simply speeding through the experiment without paying attention), >50% of trials were >5000 ms (suggesting they were not fully focusing on the experiment), or they chose the same button (left or right) on >90% of trials (suggesting they were simply pressing buttons to proceed without paying attention). Across all three criteria, this resulted in 9 participants being excluded from study 1a (1 with short latencies, 6 with long latencies, and 2 with little variation in responses), 9 participants from study 1b (3 with short latencies, 5 with long latencies, and 2 with little variation in responses), and 16 participants from study 1c (11 with short latencies, 3 with long latencies, and 5 with little variation in responses; note that numbers in parentheses do not add up to total numbers excluded because some participants met multiple exclusion criteria). This left final sample sizes of 94 for study 1a, 179 for study 1b, and 284 for study 1c. We report the results with these participants removed below, but results are similar if we instead leave these participants in.

5.2.2. Analyses

In all analyses predicting participants' choices, we conducted multilevel logistic regression models using the lme4 package (Bates, Mächler, Bolker, & Walker, 2015) in R (R Core Team, 2021), with random intercepts for subjects. Continuous predictors were centred and standardized before being entered into the model. For effects of continuous predictors and two-level categorical variables, we report the regression coefficient and z-values. For categorical predictors with more than two levels, we report the Type III Wald X^2 values, obtained using the car package (Fox & Weisberg, 2019). Follow-up pairwise comparisons are conducted using the emmeans package (Lenth, 2021), with Tukey adjustments for multiple comparisons. Model specifications for each analysis are detailed below. These studies were not preregistered. Data and analysis code for all studies in this paper are available on the Open Science Framework at https://osf.io/ud7cr/?view_only=e42e1ce67b9340b08cd1a43cffe9352.

5.2.3. What do people choose?

Before examining the main question of interest, we first wanted to check that participants were using the information in ways that were expected in order to make their decisions. We therefore ran models predicting participants' choices from the information that was presented on screen. Since self information appeared on both sides of the screen, we first computed a difference score by subtracting the self outcome that was associated with an outgroup outcome from the self outcome that was associated with an ingroup outcome. This created a "self difference" score where higher numbers mean that choosing the ingroup side yielded better outcomes for the self than choosing the outgroup side. To examine how participants were using the different pieces of information to make their choices, we then predicted choice (coded as 0 = chose outgroup side and 1 = chose ingroup side) from the self outcome difference, the ingroup outcome, and the outgroup outcome, with trials nested within subjects and random slopes modeled for each of these variables. Raw choice proportions as a function of the four outcomes are displayed in Supplementary Fig. 2.

Study 1a. In study 1a, this analysis indicated that participants were more likely to choose the ingroup side if it led to better outcomes for themselves, ($b = 0.94, z = 7.92, p < .001$), better outcomes for their ingroup, ($b = 2.91, z = 16.29, p < .001$), and worse outcomes for their outgroup ($b = 0.87, z = 5.86, p < .001$). In other words, greater positive values for the self and ingroup predicted an increased likelihood of choosing the ingroup side, whereas greater negative values for the outgroup predicted an increased likelihood of choosing the outgroup side.

Study 1b. Study 1b replicates this result, demonstrating that participants were more likely to choose the ingroup side if it led to better outcomes for themselves, ($b = 0.73, z = 12.28, p < .001$), better outcomes for their ingroup, ($b = 2.47, z = 25.73, p < .001$), and worse outcomes for their outgroup, ($b = 0.79, z = 9.18, p < .001$).

Study 1c. Replicating the first two studies, people in study 1c chose outcomes that were good for themselves ($b = 1.28, z = 11.91, p < .001$) and good for the ingroup ($b = 3.80, z = 18.42, p < .001$). However, unlike the previous studies, no main effect of outgroup reward was found ($b = -0.05, z = -0.68, p = .495$), suggesting that participants on average were neither trying to help nor hurt the outgroup.

5.2.4. How does self-other conflict affect use of information?

The main question for studies 1a-1c revolves around how the self-related information influenced participants' use of group-related information in their decisions. If people are simply weighting different pieces of information linearly to make their decision, the difference between the two self outcomes should influence choices in a graded manner. There should be nothing special about a difference of zero between the two self outcomes; rather, the greater the difference between potential self outcomes, the more likely participants should be to use information about the self outcomes in their decision.

Alternatively, if people prioritize their goals in a non-linear manner, the mere presence or absence of opportunities to help the self should be more relevant. Thus, if participants are instead weighting and de-weighting information in accordance with their goal priorities, we should see a distinct difference between the use of information when the two self outcomes are equal and when the two self outcomes are unequal, with differences between other values of self differences less pronounced.¹

To test this question, we created a categorical variable representing the absolute difference between the two self outcomes. Since each self outcome had possible values of 1, 5, or 10, this difference had possible values of 0, 4, 5, or 9. We then predicted choice (ingroup side vs. outgroup side) from the ingroup outcome, the outgroup outcome, the categorical absolute self difference, and the interaction of these three variables, with random slopes included for the ingroup and outgroup outcomes.

Study 1a. In study 1a, this analysis reveals a significant interaction of the self difference and ingroup outcome, $X^2(3) = 38.62, p < .001$, with a clear separation of whether or not the self outcomes were equal. As seen in Fig. 2, participants help the ingroup more when the self outcomes are equal (i.e., the self difference is 0) compared to when the self outcomes are unequal. This increase in the use of ingroup information is non-linear: as depicted in Fig. 3, participants' use of the ingroup information does not change significantly when the self difference is 4 or 5 compared to when it is 9 (5 vs. 9 contrast: $b = 0.16, z = 1.39, p = .505$), but they use the ingroup information far more when the self difference is

0 (0 vs. 4 contrast: $b = 0.45, z = 3.77, p < .001$). This suggests that the interaction of self- and ingroup-related information is occurring through a weighted prioritization of goals rather than a simple linear use of information. The parallel interaction for the outgroup outcome was not significant, $X^2(3) = 1.19, p = .757$, suggesting that use of the outgroup information did not depend on the self difference. To further probe this null effect, we simplified this model to compare cases where the self outcomes are equal vs. unequal, and then re-ran the model as a Bayesian multilevel model using the brms package (Bürkner, 2017). We used uniform priors across the potential parameter space from -1 to 1 (using scaled variables) and ran the model for 3000 samples. The Bayes factor for the interaction between the outgroup outcome and whether the self was equal was 18.36, indicating that the null hypothesis of no effect is >18 times more likely after the data is incorporated into the model (see Fig. 2).

Study 1b. Similarly, participants' use of ingroup information in Study 1b also increased as a function of the self outcome difference. Participants helped the ingroup more when the self outcomes were equal (see Fig. 2), and the increase in the use of ingroup information as a function of the self outcome difference was non-linear (see Fig. 3). Replicating the first study, we found a significant interaction between the self outcome difference and the ingroup outcome, $X^2(3) = 48.92, p < .001$. This interaction revealed that the use of ingroup information showed a greater increase from a self difference of 4 to a self difference of 0 ($b = 0.34, z = 4.16, p < .001$) than it did from a self difference of 9 to a self difference of 5 ($b = 0.14, z = 1.61, p = .374$). Unlike Study 1a, the interaction in this study was also significant for the outgroup, $X^2(3) = 14.67, p = .002$, such that use of the outgroup information also depended on the self outcome difference. However, pairwise comparisons suggest that the use of outgroup information is only significantly different when comparing a self difference of 0 to a self difference of 9 (0 vs. 4 contrast: $b = 0.14, z = 2.04, p = .175$; 5 vs 9 contrast: $b = 0.085, z = 1.17, p = .645$; 0 vs. 9 contrast: $b = 0.25, z = 3.73, p = .001$), and the same non-linearity is not as apparent as for the ingroup (see Fig. 3).

Study 1c. Study 1c also reveals an interaction between ingroup outcome and the self outcome difference, $X^2(3) = 249.24, p < .001$. Replicating the first two studies, we found that the use of ingroup information showed a large increase when the self outcomes were equal (self difference of 0 vs. 4: $b = 0.81, z = 12.48, p < .001$), and only a slight increase with varying levels of non-zero self difference values (self difference of 5 vs. 9: $b = 0.24, z = 3.70, p = .001$; see Figs. 2 and 3). No significant interaction was found between the outgroup outcome and the self outcome difference, $X^2(3) = 0.14, p = .987$. As in Study 1a, we therefore re-ran this model in brms (using the binary selfEqual variable in place of the categorical absolute self difference variable) to further probe this null effect. We found a Bayes factor of 36.34, suggesting that the null hypothesis of no effect is >36 times more likely after the data is incorporated into the model. Thus, it seems that the use of ingroup-related information depended on the presence of self goals, but the use of outgroup-related information did not. Throughout these studies, this asymmetry in tradeoffs between self and ingroup compared to self and outgroup holds even when examining these effects separately based on the participant's overall orientation towards the outgroup (see Supplementary Materials). Further, the same asymmetry in the use of outcomes is not present when we reverse the roles of the self and group information (see Supplementary Fig. 4), suggesting that this pattern is unique to the effect of self information on group-oriented behaviour.

6. Study 2

Study 2 was very similar to Studies 1a-1c, with two notable differences designed to address potential limitations in the first set of studies. Studies 1a-1c asked participants to respond as though all choices they were making translated into real money; however, these choices were purely hypothetical and not directly tied to their monetary outcomes. To ensure that these results replicate when participants' choices are real

¹ We conducted simulations to confirm that this pattern should occur only under non-linear prioritization of self-related information. Specifically, we simulated data under a linear relationship of ingroup information and the self outcome difference and observed the resulting patterns from analysing the data using a categorical self outcome difference variable. Results are presented in Supplementary Figure 3 and confirm that the pattern does not occur under a linear relationship between these variables.




Fig. 2. Interaction of group outcomes with whether or not the self is equal in studies 1a, 1b, and 1c. For visualization purposes, we re-ran the models described in the text using the binary variable indicating whether or not the self outcomes were equal as a moderator. Panel A depicts the results of a Bayesian multi-level model run on all three studies combined, showing that ingroup helping increases when the self outcomes are equal, whereas behaviour towards the outgroup is largely unaffected. Panel B contains the Bayesian estimates and 95% credible intervals for each study separately, as well as all three studies combined.




Fig. 3. Effect of group outcomes on choices at different levels of self information. For the ingroup (left panel), there is a non-linear increase in use of the group information as a function of the self information. When the difference between the two self outcomes is 0 (i.e., no progress can be made towards self-related goals), participants show a large increase in the actions they take for the ingroup. A similar but weaker and inconsistent pattern appears to be present for the outgroup.

rather than hypothetical, we incentivized participants' choices in Study 2 with real money, such that all choices they made both for themselves and for others determined how much bonus money each person received. For example, a given participant who was randomly assigned to the red team would receive a final bonus amount that was proportional to the total points the participant received for themselves, the total points the other players on the red team gained for the ingroup, and the total points the players on the blue team gained (or lost) for the outgroup.

Studies 1a-1c revealed that the interdependency between self-related goals and ingroup-related goals was stronger than the interdependency between self-related and outgroup-related goals. However, it is possible that this difference simply arose out of the need to balance multiple types of information at once. Since both ingroup- and outgroup-related information was present in each trial, it may have been difficult for participants to shift their attention to *both* the ingroup and the outgroup when the self outcomes were equal. In other words, it is possible that participants would have shifted their attention more to the outgroup if no ingroup goal were present. Thus, in Study 2, we separated ingroup- and outgroup-related information across blocks, so that in a single block participants were either making decisions for the self and ingroup or for

the self and outgroup, and not all three simultaneously.

6.1. Methods

6.1.1. Participants

As this study was intended to additionally serve as a pilot study for a future fMRI study, we deliberately restricted our sample size to that feasible for an fMRI study to examine whether we could still reliably detect these (within-subject) effects. We therefore recruited a sample of 40 participants and conducted a sensitivity analysis to determine our ability to reliably detect our effect of interest using a z-test. Similarly to the first three studies, this analysis indicated that we had 80% power to detect interaction slope estimates as small as -0.11 .

40 undergraduate students (29 East/Southeast Asian, 2 White, 2 South Asian, 1 Black, 4 other races/ethnicities, 2 did not report race/ethnicity; 28 female, 10 male, 2 did not report gender; mean age = 20.7, SD age = 2.9) were recruited from the University of Toronto community and paid for participating in the study. Participants were given a base payment of \$10 for participating and could earn an additional bonus amount of up to \$5 depending on their and others' performance in the task.

6.1.2. Procedure

The task used in this study was similar to the previous studies, with three main changes. First, ingroup and outgroup information was separated across blocks, such that participants were either making decisions for the self and the ingroup or the self and the outgroup, not all at once. Within these blocks, potential outcomes for the self were +1, +5, and +10 points, and potential outcomes for a random group member were -5 , -1 , +1, +5, +10, and +15 points. Second, we used real money to incentivize participants' choices - participants were informed that their final payment from the study would depend on both their own and others' decisions. Specifically, their bonus money would be a function of the points they gained for themselves, the points people on their own team gained for their team, and the points the people on the other team gained or lost for them. Participants were not told the exact translation of points to bonus money, but were told that they could earn up to a maximum of \$5 as a bonus. Finally, as this study also served as a pilot for an fMRI study, the structure of the trials and blocks was changed. Specifically, participants completed a block of 6 practice trials followed by 4 "runs" of trials. Each run was divided into 8 blocks, 4 of which involved choices for the self and ingroup, and the other 4 involving choices for the self and outgroup. Participants had 5 s to make their response on each

trial, after which the trial would time out.

After completing the task, participants filled out questionnaires about their demographics and their level of identification with their assigned group (Leach et al., 2008). The group identification questionnaire consists of 14 items measuring the degree to which someone identifies with a given group. Three of these items did not make sense in a minimal group context and were thus removed, leaving 11 items which participants responded to on a 1 (strongly disagree) to 7 (strongly agree) scale. Analyses of this questionnaire are presented in the Supplementary Materials. The same exclusion criteria as in Studies 1a-1c were applied in Study 2, but no subjects met these criteria and needed to be excluded.

6.2. Results

6.2.1. What do people choose?

Since each trial in this study presented participants with two options for the self and two options for either the ingroup or the outgroup, we first computed difference scores representing these variables. Specifically, we computed the difference between the two self options and the difference between the two group options. We then predicted choice (coded as choosing the left vs. right side) from the two difference scores, which were interacted with a variable representing which group was in the current block (the ingroup or the outgroup). Random slopes were also included for the two difference scores. This analysis indicated that people chose outcomes that were good for themselves, $b = 2.65$, $z = 7.90$, $p < .001$, and interestingly this effect was slightly greater in ingroup blocks than outgroup blocks, $b = 0.16$, $z = 3.53$, $p < .001$. We found an overall main effect of the group difference, such that participants in general chose things that were good for the groups, $b = 0.68$, $z = 3.09$, $p = .002$. However, this was qualified by a large two-way interaction between the group difference and the group (ingroup vs outgroup), $b = 1.62$, $z = 32.77$, $p < .001$, indicating that participants chose positive outcomes for the ingroup ($b = 2.30$, $z = 10.03$, $p < .001$) but negative outcomes for the outgroup ($b = -0.93$, $z = -4.16$, $p < .001$).

6.2.2. How does self-other conflict affect use of information?

To examine how self-other conflict influenced decisions in this task, we predicted choice (coded as left vs. right) from the difference between group options, the group (ingroup vs. outgroup), the categorical absolute difference between the potential self outcomes, and the interaction of these variables. Trials were nested within subjects, and random slopes were modeled for the group outcome difference. This analysis indicated that the use of group information depended on the absolute self difference, $X^2(3) = 45.96$, $p < .001$. However, this was further qualified by a three-way interaction between the group outcome difference, the group (ingroup or outgroup), and the absolute self difference, $X^2(3) = 91.69$, $p < .001$. As seen in Fig. 3, participants once again use the ingroup information much more when the self outcome difference is 0 compared to when it is anything else (0 vs 4 contrast: $b = 1.25$, $z = 8.24$, $p < .001$; 5 vs. 9 contrast: $b = 0.24$, $z = 1.90$, $p = .227$). The parallel comparisons for the outgroup are not significant, (0 vs 4 contrast: $b = -0.20$, $z = -1.81$, $p = .268$; 5 vs. 9 contrast: $b = -0.089$, $z = 0.75$, $p = .878$), suggesting the use of outgroup information does not depend as heavily on the self information.

6.3. Interim summary

Studies 1 and 2 examined which tradeoffs are made when balancing multiple proximal goals in social situations, finding that ingroup-oriented actions are more likely to trade off with self-oriented actions than are outgroup-oriented actions. While people had goals for both the ingroup and the outgroup, progress towards ingroup goals was more dependent on the self goals, pointing to the potentially greater interdependent nature of the self and the ingroup. Further, these studies

found that people weigh goal-relevant information non-linearly, upweighting the ingroup goal especially when no progress could be made towards the self goal, in a manner indicative of sequential rather than concurrent goal pursuit. These findings replicated multiple times throughout Studies 1a-c and 2, including when real payoffs were used and ingroup and outgroup information was separated across blocks.

7. Study 3

The results of our behavioural studies provided evidence that goals for the self and groups are pursued sequentially rather than concurrently, with ingroup-related goals upweighted when self-related goals are blocked. In study 3, we turn to examining the mechanisms of this upweighting, asking *how* people shift to prioritizing group-related goals when they cannot benefit themselves. Using eye-tracking, we examine how tradeoffs between self- and other-related goals unfold in time, investigating the dynamics of how people select and integrate information under multiple goals.

Two non-mutually exclusive possibilities exist for how competing goals may affect behaviour, outlined in Fig. 4. First, competing goals may shape people's attention to different pieces of information, such that goal competition is resolved through attentional mechanisms. For example, if a self-motivated perceiver is faced with a choice to benefit only themselves or to benefit others in their group, they may focus more heavily on the stimuli and potential behaviours that are relevant to helping themselves, ignoring those relevant for helping others. Second, these competing goals may influence the way they use the information that has been attended to, with goal competition resolved after information is gathered and when that information is instead being weighed for decision-making. For example, the self-motivated perceiver might pay equal attention to the stimuli relevant to helping themselves and others, but simply use that information differently, choosing to utilize the self-relevant information more than the other-relevant information. This distinction could have important implications for interventions that aim to increase prosociality, especially towards outgroup members. If the lack of other-regarding behaviour in a particular situation is simply caused by inattention to others, drawing attention to potential recipients in need should be enough to increase altruism. However, if during goal competition people are attending to others but simply neglecting to use that information in a prosocial manner, drawing attention to potential recipients would not be beneficial, and may even increase harmful behaviours. Here, we aim to gain some insight into these questions by examining the dynamics of how we select and integrate information under multiple goals, testing whether goal tradeoffs shape people's attention to information and/or influence the way people use the




Fig. 4. Potential roles of goal conflict in influencing goal-directed behaviour. Two possible mechanisms are depicted for the role of goal conflict on goal-directed behaviour. First, when goals conflict, people may modify their search for information, looking more to goal-relevant stimuli (pathway depicted in blue). Second, goal conflict may influence the relationship between eye gaze and the use of information in decision-making (pathway depicted in green), such that attending to information has a greater effect on the use of that information when conflicting goals are not present. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

information that is attended to.

7.1. Methods

7.1.1. Participants

83 undergraduate students from the University of Toronto were recruited to participate in the study. Due to technical issues, the experiment was not run or data was not saved for 10 participants, leaving a total of 73 participants for analysis (47 female, 25 male, 1 did not report gender; 45 East/Southeast Asian, 11 South Asian, 8 White, 5 Middle Eastern, 2 mixed race, 1 other race, 1 did not report race; mean age 19.8, SD age 1.3). Participants received partial course credit for completing the study, and also had the opportunity to gain up to \$10 in bonus money depending on the choices made by them and other participants in the task.

As in the previous studies, we conducted a sensitivity analysis to determine the smallest behavioural effect of interest we could reliably detect using a *z*-test. This analysis revealed that we had 80% power to detect interaction slope estimates as small as -0.07 .

7.1.2. Eye-tracking apparatus

Eye movement data was collected using an EyeLink 1000 Plus with a sampling rate of 500 Hz, with participants' heads resting on a desktop mount to minimize movement. Stimuli were presented on a 24-in. display screen (resolution 1920×1080 , refresh rate 60 Hz), with participants seated so their heads rested in a chin rest approximately 80 cm from the screen.

7.1.3. Task

The task used in this study was very similar to that of the earlier studies: participants were randomly assigned to the red team or the blue team and informed that they would play a game in which they made decisions for themselves and others while we track their eye movements. In this study, instead of presenting the possible points for the entire team in addition to each team member (as we did in Study 1), we instead randomly presented one cartoon member from the team and presented the outcome for that person specifically. Possible values for the self outcomes were +1, +5, and +10, while possible values for the group outcomes were -5 , -1 , +1, and +5. If a response was not given within 10s of the trial appearing, the trial timed out and the task proceeded to the next trial.

The decisions that participants made in the task directly affected the bonus money earned by themselves and by the other participants in the study. Specifically, participants were informed that their final bonus money would be determined by the points they got for themselves, the points the other members of their team got for the ingroup, and the points the members of the other team got for the outgroup. Similarly, the decisions the participants made would affect the bonus money of themselves, their team members, and the other team's members. Thus, all decisions in the study translated into real monetary outcomes both for oneself and for others. Participants were not told the exact translation of points to dollars, but were informed that they could gain up to \$10 in bonus money.

7.1.4. Procedure

As in the earlier studies, after reading the task instructions and being randomly assigned to either the red team or the blue team, participants were presented with a set of avatars varying along race and gender and were asked to choose the one that most resembled them to represent them in the task. This avatar was then used to represent the self on each trial, with the ingroup and outgroup members represented by random avatars that changed on each trial.

Before beginning the task, participants underwent a calibration and validation procedure, in which 9 dots sequentially appeared on screen and participants were instructed to fixate on each one. This procedure ensures that the eye-tracker is accurately able to map computed fixation

positions onto the actual targets on screen.

Upon completing calibration and validation, participants completed a set of 5 practice trials before starting the main experiment. The experiment was divided into three blocks of 100 trials each, with trials separated by fixations. Calibration and validation were performed at the beginning of each block.

Once the final block was completed, participants completed a series of questionnaires. This included questionnaires assessing general demographics, the social dominance orientation scale used in Study 1, and the group identification scale used in Study 2. Analyses of these questionnaires are presented in the Supplementary Materials. Some participants also completed a scale measuring their social value orientation (Murphy & Ackermann, 2013); however, due to a technical error, data on this scale was only collected for a subset of participants and thus is not discussed further. The same exclusion criteria as in Studies 1a-1c were applied in Study 3, but no subjects met these criteria and needed to be excluded.

7.1.5. Processing eye-tracking data

To assist in analysis of the eye-tracking data, we defined a series of non-overlapping visual interest areas. This included four main interest areas encapsulating the text of each of the four possible outcomes (each with a pixel area of 19,950). To supplement this analysis, we also defined four additional interest areas for the four cartoon pictures that were shown on screen (each with a pixel area of 38,000) and an interest area for the reminder text informing participants of the key mapping at the top of the screen (with a pixel area of 41,310). The text interest areas contained the information necessary for participants to make their decisions, and thus are the main focus of our analysis. On average, 9.9% of all fixations were made to the picture interest areas; these fixations are not of primary interest here and are not discussed further.

To clean the eye-tracking data, we examined the proportion of samples in each trial where no eye-tracking data was recorded, such as when the participant was blinking. We then dropped any trials whose trackloss proportion was >0.25 , resulting in 346 trials across all subjects being removed (including an average of 3.8% of trials per participant and a maximum of 16.5% of trials for one participant). This left an average of 292 trials (median 303) per participant for analysis.

Finally, we examined the distributions of participants' response times. The average time to make a decision across all trials was 3.2 s, with a strong right skew. In the analyses that follow in which we collapse across the time within the trial, we include the entire trial length in our analysis. In analyses where time within the trial is explicitly analyzed (i. e., the divergence analysis and GAMs), however, we exclude times >6000 ms, which account for only 3.03% of the overall data.

7.2. Results

7.2.1. Preliminary effects

What do people choose? We first examined how participants used the information presented to them on each trial to make their choices. To do so, we ran a multilevel logistic regression model predicting choice (coded as 1 = chose ingroup side, 0 = chose outgroup side) from the ingroup outcome, the outgroup outcome, and the difference between the two self outcomes, with these variables included as both fixed effects and random effects. Replicating the previous studies, this analysis indicated that participants chose outcomes that were positive for the ingroup, $b = 2.18$, $z = 15.27$, $p < .001$, positive for the self, $b = 2.32$, $z = 10.39$, $p < .001$, and negative for the outgroup, $b = 0.95$, $z = 6.62$, $p < .001$.

How does self-other conflict affect choice? To examine how self-other conflict affected participants' choices in this task, we predicted choice from the ingroup outcome, the outgroup outcome, and the categorical absolute difference between the potential self outcomes. Random slopes were once again modeled for the ingroup and outgroup outcomes, and trials were nested within subject. Replicating earlier

analyses, the use of ingroup outcomes varied as a function of the absolute self difference, $\chi^2(3) = 486.18, p < .001$, with participants helping the ingroup more when their self-related goals were not relevant. As seen in Fig. 3, participants use the group information more when the self outcome difference is 0 compared to when it is anything else (0 vs 4 contrast: $b = 0.67, z = 11.65, p < .001$; 5 vs. 9 contrast: $b = 0.36, z = 6.85, p < .001$). A parallel but smaller interaction was found for the outgroup outcomes, with participants hurting the outgroup more when the two self outcomes were equal, $\chi^2(3) = 81.26, p < .001$ (0 vs. 4 contrast: $b = 0.24, z = 4.65, p < .001$; 5 vs. 9 contrast: $b = 0.15, z = 2.89, p = .020$; see Fig. 3).

What do people look at? The outcomes that participants look at first on each trial can provide some indication of their informational priorities. We therefore examined participants' first fixation on each trial to see which information they were looking at, collapsing across fixations to text and images. Since the position of information varied between subjects and had a large effect on participants' first fixation (with participants typically looking first at the information just to the left of the starting fixation position), we controlled for position in this analysis. Specifically, we examined the average proportion of first fixations to each interest area for each of the four position configurations, and then averaged across these four results to get an overall estimate of which information participants tended to look at first. Overall, participants most often looked at the ingroup side first, with first fixations to the ingroup on 29.3% of trials, the self on the ingroup side on 28.6% of trials, the outgroup on 24.1% of trials, and the self on the outgroup side on 19.7% of trials.

After establishing where participants' first fixations tend to be, we can also look at where people tend to spend the most time looking overall. We use a poisson multilevel model predicting fixation counts from the interest area, with subject as a random intercept, to analyse the overall number of fixations to each interest area. Overall, the number of fixations differed by interest area, $\chi^2(3) = 3100.40, p < .001$, with participants making an average of 1.94 fixations per trial to the ingroup outcomes, 1.84 fixations to the outgroup outcomes, and 1.42 and 1.40 fixations to the self outcomes on the ingroup and outgroup sides respectively.

This pattern replicates if we instead conduct an overall window analysis, looking at the proportion of time participants spend looking at each interest area within a trial. As the gaze proportions are bounded between 0 and 1, we first applied a corrected logit transformation to the proportions using the eyetrackingR package (Dink & Ferguson, 2018), which we used as our dependent variable. Specifically, predicting the transformed gaze proportion to each interest area from the interest area, with subject as a random effect, suggests that people spend the greatest proportion of time looking at the ingroup (mean proportion = 0.30), followed by the outgroup (mean proportion = 0.29) and then the self on the ingroup side (mean proportion = 0.26) and outgroup side (mean proportion = 0.23) information, $\chi^2(3) = 1740.71, p < .001$.

How does attention affect choice? As a final preliminary analysis, we can combine the analysis of choice and gaze data to examine how participants' attention to the different pieces of information affected the use of that information in their decisions. To calculate a series of gaze variables, we took the proportion of samples to each interest area within a trial and then calculated a corrected logit transformation of these variables using the eyetrackingR package. We then ran a multilevel logistic regression model predicting choice (1 = chose ingroup side; 0 = chose outgroup side) from the four potential outcomes interacted with the participant's gaze to each outcome in that trial, with a random intercept modeled for subject and random slopes modeled for the group outcomes. This analysis indicated that the more a participant looked at a particular piece of information, the more it affected their choices. This occurred in the positive direction for the self (self on ingroup side: $b = 0.068, z = 8.35, p < .001$; self on outgroup side: $b = -0.075, z = -8.38, p < .001$), and the ingroup ($b = 0.11, z = 7.97, p < .001$), such that the more a participant looks at the self or ingroup information, the more

likely they are to make choices that are good for the self or ingroup. In contrast, the more the participant looked at the outgroup information, the more likely they were to choose negative outcomes for the outgroup, $b = 0.053, z = 4.05, p < .001$.

This effect could be occurring at the subject level, such that someone who tends to look more at a given piece of information also tends to use that information more in their decision, and/or at the trial level, such that if on a particular trial a participant looks more at a particular piece of information, they will be more likely to use that information. To examine these two possibilities, we separated each gaze variable into two orthogonal variables: the mean gaze for each person centred between-subjects, and the trial-level gaze centred within-subject. We then re-ran the previous model with these new gaze variables in place of the overall gaze variables used previously. This analysis revealed that the effects of gaze largely held at both the subject-level and the trial-level. Specifically, these effects held for the ingroup (although only marginal at the between-subject level; subject-level: $b = 0.24, z = 1.73, p = .084$; trial-level: $b = 0.19, z = 8.00, p < .001$), the outgroup (subject-level: $b = 0.33, z = 2.67, p = .008$; trial-level: $b = 0.082, z = 3.69, p < .001$), the self on the ingroup side (subject-level: $b = 0.15, z = 7.00, p < .001$; trial-level: $b = 0.094, z = 4.28, p < .001$), and the self on the outgroup side (subject-level: $b = -0.21, z = -9.53, p < .001$; trial-level: $b = -0.06, z = -2.66, p = .008$). Thus, this provides evidence of both individual differences in gaze and choice strategy as well as intra-individual variation across trials.

7.2.2. How does self-other conflict affect information search?

Our first set of hypotheses concerns whether self-other conflict affects the kinds of information participants look for to make their decision. In other words, we examine how the presence of self-related goals might shape participants' explicit attention to different pieces of information.

Window analysis. We can first conduct a window analysis on the entire length of the trial, asking how self-related goals shape attention to each interest area. Specifically, we predict the corrected log of the participants' gaze proportion from the interest area, a variable indicating whether or not the self outcomes were equal, and the interaction of the two, with a random intercept for subject and a random slope for the selfEqual variable. In addition to main effects for each variable (selfEqual: $\chi^2(1) = 23.32, p < .001$; interest area: $\chi^2(3) = 32.17, p < .001$), this analysis yields a significant interaction, $\chi^2(3) = 176.87, p < .001$, suggesting that when the two self outcomes are equal, participants spend less time looking at the self outcomes (self on ingroup side: $b = 0.15, z = 7.96, p < .001$; self on outgroup side: $b = 0.17, z = 9.48, p < .001$), and more time looking at the group outcomes, (ingroup: $b = -0.086, z = -5.12, p < .001$; outgroup: $b = -0.061, z = -3.57, p < .001$).

Divergence analysis. The window analysis indicates that the relevance or irrelevance of self-related goals influences people's overall attention to the different pieces of information. This suggests that some kind of attentional prioritization is taking place, but to better understand when this prioritization emerges we need to examine the timecourse of these effects within the trial. In particular, we want to know when in the trial self-other conflict had an effect on fixations to the ingroup and the outgroup. We therefore ran a divergence analysis to discover the onset of the effect, as well as how long the effect lasted for within the trial.

To account for clustering of time bins while controlling the Type I error rate, we ran a bootstrapped cluster-based permutation analysis using the eyetrackingR package. This analysis involves repeatedly shuffling time-bins within the data and calculating t-values of time-clusters (i.e., groups of adjacent time-bins), saving the biggest t-value on each iteration to create a null distribution. Significant clusters from the non-shuffled data are then compared to this null distribution to determine statistical significance. We ran two such divergence analyses, with 2000 samples each, to examine when in the trial self-other goal conflict affected gaze to the ingroup and outgroup respectively.

Analysing gaze to the ingroup revealed a significant effect of whether the self was equal from 600 ms to 1800 ms, $t = -36.17$, $p < .001$, such that during this time period people looked more at the ingroup when the self outcomes were equal compared to when they were unequal. A parallel analysis on the outgroup suggested that the effect on gaze to the outgroup occurs from 1700 ms to 2300 ms, $t = -14.56$, $p = .009$ (see Fig. 5a).

GAMs. After establishing that the effect of self-other conflict varies over the course of a single decision, we next wanted to examine whether this temporal prioritization shifts over the course of many decisions. Participants completed 300 trials in our experiment, which we intentionally made longer than average to examine how prioritization might change with fatigue or boredom. We can therefore examine how this time-varying effect of self-other conflict shifts over the course of the entire experiment.

To examine this question, we ran a series of generalized additive models (GAMs) using the *bam* function from the *mgcv* package (Wood, 2011). In each model, we predicted the proportion of samples to a particular interest area from the time within trial, the trial number, and a variable indicating whether or not the self outcomes were equal. All models included an AR1 error model to account for autocorrelation in adjacent time bins, with the ρ parameter estimated from an initial fit using the *itsadug* package (van Rij, Wieling, Baayen, & van Rijn, 2020) and set to a value of 0.82. Our main models of interest included a parametric term for whether or not the self outcomes were equal and a tensor product smooth representing the interaction of time within trial, trial number, and whether the self was equal. The random effects structure for these models included a non-linear factor smooth over time for subject and an ordered factor difference smooth (as described in Wieling, 2018) for whether or not the self was equal. Model comparison using a chi-square test on the fREML scores suggested that this random effect structure is preferred over a model with no random effect for whether or not the self was equal (and only a random smooth over time for subject; ingroup: $X^2(2.00) = 159.34$, $p < .001$; outgroup: $X^2(2.00) = 137.17$, $p < .001$), as well as a model in which the effect of selfEqual can vary by subject but not over time (i.e., a random slope rather than a random smooth; ingroup: $X(1.00)^2 = 146.22$, $p < .001$; outgroup: $X(1.00)^2 = 109.29$, $p < .001$).

To test whether the interaction of selfEqual with time and trial improved the fit of the model, we re-specified the model with selfEqual as an ordered factor (Wieling, 2018). The model therefore includes one term for the smooth interaction of time and trial, and another term representing the ordered factor difference smooth for selfEqual. When predicting gaze to the ingroup, this ordered factor difference smooth is significant (see Table 1), suggesting that the addition of selfEqual to the model does improve the fit. When predicting gaze to the outgroup, however, this term does not explain a significant amount of variance

(see Table 1).²

As we found somewhat more evidence for this three-way interaction for the ingroup than the outgroup, we then return to our original model specification to further probe the interaction for the ingroup. The full results of these models are presented in Table 1. Examining the edf values for the smooth terms in the model suggests that there is evidence of non-linearity in these effects, with edf values well above 1. This effect is visualized for the ingroup in Fig. 5b, which suggests that as trials proceed, the effect of selfEqual on looking to the ingroup gets earlier and earlier in time. At around trial 50, this effect is significant from 1480 ms to 3685 ms. After the first block of 100 trials, this effect seems to shift earlier and by trial 250 is significant from 944 ms to 1957 ms.

7.2.3. How does self-other conflict affect information use?

The second set of hypotheses we aim to investigate concerns the role of self-other conflict in the use of information. In other words, does the presence or absence of self-related goals affect not just the type and timing of information that people acquire, but also the way they use that information to make their decisions? To answer this question, we must examine differences not just in participants' attention to information, but in the way that attention translates into behaviour.

We therefore ran a model predicting participants' choices on each trial from the group outcomes, the participants' gaze to the group outcomes, and a variable indicating whether or not the self outcomes were equal. Trials were nested within subject, and random slopes were included for the two group outcomes. This analysis revealed a three-way interaction between the ingroup outcome, gaze to the ingroup, and the selfEqual variable, $b = -0.035$, $z = -3.21$, $p = .001$. This result suggests that the amount the participant looked at the ingroup outcome had a greater effect on their decision (i.e., use of the ingroup outcome) when the self outcomes were equal (see Fig. 6). In other words, someone who spent a greater proportion of time looking at the ingroup outcome on a particular trial was more likely to make a choice that was good for the ingroup, but this was especially the case when the self outcomes were equal and the self-related goal was therefore irrelevant. The parallel interaction for the outgroup was not significant, $b = -0.015$, $z = -1.62$, $p = .105$.

To examine whether this effect was occurring primarily at the subject level or the trial level, we then separated each group gaze into two variables: the participant's overall average gaze to the group, centred across participants, and the trial-level deviation from the participant's average gaze. In this analysis, the three-way interactions between the group outcomes, gaze to the groups, and whether or not the self was equal were significant at the subject level but not the trial level. Significant interactions were found at the subject-level for both the ingroup, $b = -0.094$, $z = -4.54$, $p < .001$, and the outgroup, $b = -0.071$, $z = -3.85$, $p < .001$, although we mainly focus on

² Our dependent variable (proportion of samples fixated on a given interest area) consists of mostly 0 s and 1 s with about 5.6% of values falling somewhere in between these extremes. While the models reported above are gaussian models with identity link functions, the distribution of our dependent variable resulted in non-normal distributions of residuals, which is not solved through transformations of the dependent variable. To ensure the robustness of these results, we therefore excluded any data points that fell between 0 and 1 and re-ran these analyses using binomial models. This led to similar (and seemingly more robust) results, with a significant smooth interaction of selfEqual with time and trial when predicting gaze to the ingroup, $edf = 7.83$, $Ref.edf = 8.99$, $X^2 = 24.6$, $p = .004$, but not when predicting gaze to the outgroup, $edf = 3.05$, $Ref.edf = 3.094$, $X^2 = 1.23$, $p = .489$. The results of this binomial model should also be interpreted with caution since, in addition to requiring us to drop some data points, accounting for autocorrelation within these models requires us to use the "discrete" setting in the *bam* function, which discretizes the continuous data and, in our case, produces slightly different results even when used with a gaussian link function. For these reasons, we primarily report the results of the gaussian model here.




Fig. 5. Effect of self-other goal conflict on gaze towards ingroup and outgroup over time. Panel A depicts the results of a divergence analysis conducted on gaze to the ingroup and outgroup, indicating that the presence of self-related goals affects gaze towards the ingroup earlier and the effect lasts for longer than for the outgroup. Panel B depicts how this effect on gaze to the ingroup changes over the course of the study, with the effect on looking to the ingroup getting earlier and shorter across trials. Shaded areas indicate regions of time where the effect of interest is statistically significant.

Table 1

Results of generalized additive models predicting gaze to ingroup and outgroup outcomes. Proportion of samples to each interest area was predicted from the time within trial, the trial number, and a variable indicating whether or not the self outcomes were equal. An AR1 error model was used to account for autocorrelation in adjacent time bins. Edf = effective degrees of freedom; Ref.df = reference number of degrees of freedom used for hypothesis testing.

Parametric coefficients	Ingroup				Outgroup			
	Estimate	Std Error	t	p	Estimate	Std Error	t	p
Intercept	0.246	0.012	20.803	<0.001***	0.246	0.010	24.502	<0.001***
selfEqual	0.015	0.004	3.695	<0.001***	−0.004	0.004	−0.942	0.346
Smooth terms								
	edf	Ref.df	F	P	edf	Ref.df	F	p
te(Time,Trial)	21.677	22.87	26.417	<0.001***	19.26	21.19	12.531	<0.001***
te(Time,Trial,selfEqualO)	8.712	10.40	2.254	0.0121*	10.57	12.77	1.495	0.0932
s(Time,subject)	558.348	647.00	31.868	<0.001***	543.57	647.00	27.067	<0.001***
s(Time,subject,selfEqualO)	318.846	647.00	1.422	<0.001***	300.09	647.00	1.267	<0.001***




Fig. 6. Effect of self-other goal conflict on the use of ingroup information. When the self outcomes are equal (orange line), the relationship between looking at the ingroup information and using that information for decision-making becomes stronger than when the self outcomes are unequal (green line). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

interpretation of the ingroup effect since the initial three-way interaction was not significant for the outgroup. In contrast, the three-way interactions with gaze at the trial-level (removing any subject effects) were not significant for the ingroup, $b = -0.012$, $z = -0.48$, $p = .628$, or the outgroup, $b = 0.028$, $z = 1.31$, $p = .191$. Thus, these results suggest that individual differences in the degree to which participants tend to look at the ingroup had a greater effect on their use of the ingroup information when the self outcomes were equal. In other words, the absence of self-related goals allowed these individual differences in gaze to more strongly govern participants' choices.

8. General discussion

The current work explored goal tradeoffs when making decisions for the self, an ingroup, and an outgroup. A series of initial behavioural studies provided two key insights about these tradeoffs. First, we found evidence for sequential goal pursuit: people weigh self-related information non-linearly, with a large increase in other-oriented actions when no opportunity to help themselves exists. Second, tradeoffs between self- and ingroup-related goals were stronger than those between self- and outgroup-related goals, with a larger increase in the use of ingroup information than outgroup information when the self outcomes were equal. These results suggest that people tend to pursue their predominant self-related goals when any progress can be made towards them; only when such progress is impossible will they more strongly upweight other goals, especially goals to help ingroup members. We

note that processing was not wholly sequential; even under the presence of self-related goals, some positive weighting of ingroup outcomes was still present. Critically, however, this weighting increased sharply when self-related goals were absent, providing evidence of non-linearity. We then turned to eye-tracking to further explore the mechanisms of these tradeoffs, finding that competing goals shaped both participants' attention to goal-related information and the way they used the information that they had gathered. Thus, these results suggest that when self-related goals cannot be furthered, people both actively look for opportunities to help their ingroup and make better use of the information they already have with which to help their ingroup. These results replicated across a variety of task variations, including when points translated into real bonus money and when outcomes for the ingroup and outgroup were separated across blocks.

The results support the notion put forth by Gaertner and colleagues (Gaertner et al., 1999, 2002, 2012) of the motivational primacy of the individual self. Bolstering the notion of a motivational system in which self-related information is primary, participants in our studies sequentially pursued goals for themselves and their groups, with pursuit of group goals remaining much lower when the self goal could be pursued. This research sheds light on some of the processes through which the motivational primacy of the self may be enacted, revealing the attentional mechanisms of goal prioritization. Specifically, participants seemed to direct their attention and decisions towards the primary goal when it could be pursued; when that goal was blocked, they both shifted their attention to group-related information and changed the way they used that information to make their decisions. Thus, the results help to explain how the motivational primacy of the self influences decisions in multi-goal environments.

The findings further suggest that simultaneously examining drives towards self-interest and intergroup bias can provide a more externally valid context for understanding social behaviour. In particular, intergroup biases may rarely be as simple as helping the ingroup and hurting the outgroup; rather, in at least some cases it can take powerful social motivations to overcome self-interest and choose to instead prioritize the ingroup.

In revealing processes at the intersection of self-interest and ingroup bias, the results have a number of implications for efforts at changing behaviours. In some cases, these results suggest that harmful intergroup behaviours could potentially be reduced by highlighting when the choice to unfairly privilege the ingroup is not in the actor's self-interest (Stagnaro et al., 2018). For example, increasing interracial working class solidarity could highlight the fact that racism is often not in the interests of White workers (e.g., see Brueggemann & Boswell, 1998; Mahoney, 1999). In such cases, it is possible that increasing self-interest could actually reduce harmful ingroup-focused behaviours. Further, the results suggest that there could sometimes be negative consequences to trying to generally reduce self-interest without regard to the potential targets of other-focused behaviour, since doing so could lead to selective increased investment in the ingroup at the expense of the outgroup. Indeed, some research shows that those who have higher tendencies towards general prosociality are also more likely to demonstrate parochial forms of altruism in which they help their ingroup members but not outgroup members (e.g., Thielmann & Böhm, 2016). These findings thus support the notion put forth by Gaertner, Sedikides, Luke, and Iuzzini (2008) that a complete silencing of self-interest would be detrimental for intergroup relations. Finally, the attentional mechanisms revealed in study 3 suggest that simply drawing people's attention to outgroup members can in some cases be successful. If people neglect outgroups in part because they lack sufficient attentional resources to prioritize them, simply increasing the salience of outgroup members in need could sometimes increase prosociality towards them (though in more competitive contexts could potentially increase outgroup-directed harm). However, overt attention did not fully mediate these effects, with additional effects on the use of attended information suggesting that such interventions may be only partially successful.

8.1. Implications for goal structure

Taken together, these results help shed light on the cognitive structuring of interrelated goals in intergroup decision-making, with a few possible implications for goal structuring. Under sequential goal pursuit, secondary goals can be "shielded" from pursuit by the primary goal, which receives the bulk of resources (Shah et al., 2002). This conceptualization suggests that self-related goals in these studies shielded the group-related goals from pursuit, such that the latter were only pursued when progress towards self-related goals was blocked. When the self-related outcomes were equal and the high-magnitude goal of helping the self was made irrelevant, this shielding was removed and the next most important goal of helping the ingroup was upweighted.

However, rather than conceptualizing goals for the self, ingroup, and outgroup as three wholly separate high-level goals, the stronger tradeoffs found between the self and the ingroup may instead suggest a nested hierarchical structure for these goals. It is possible that people view ingroup helping as beneficial to their self-interest in the long run, such that the goal of helping the ingroup represents a nested goal within the goal of helping oneself, with the goal of hurting the outgroup remaining more separate. In other words, helping the ingroup may be seen as a means to helping the self, producing a state of "equifinality" (Fernandez & Kruglanski, 2018; Kruglanski et al., 2018; Kruglanski, Chernikova, Babush, Dugas, & Schumpe, 2015) in which multiple means for ultimately benefiting the self (i.e., by helping directly or by helping indirectly through the ingroup) are present. As these means are substitutable, when the goal to directly help the self is blocked (because the self-related outcomes are equal), people may instead choose to help the ingroup as a means of ultimately benefiting the self.

While the current work cannot definitively arbitrate between these possibilities, the notion of ingroup helping as a primarily self-focused phenomenon has received support from a variety of related literatures. Motivations to help oneself and one's ingroup are often highly overlapping; in fact, people commonly believe that cooperating with an ingroup is in their self-interest even when it is not (Baron, 1997, 2001). De Cremer and Van Vugt (1999) suggest that ingroup identification transforms self-interested motivations into group-interested ones, such that outcomes for the self and group become almost motivationally interchangeable. The cognitive representation of self and ingroup also tends to be blurry, especially under conditions of high group identification. Swann and colleagues (Swann et al., 2009; Swann, Jetten, Gómez, Whitehouse, and Bastian, 2012) propose the idea of self-ingroup identity fusion, in which one's personal and group-based identities fuse to become functionally equivalent. This kind of ingroup fusion predicts ingroup-oriented behaviours, with those with more highly fused identities more likely to sacrifice their own benefit to help ingroup members (Purzycki & Lang, 2019). Further, when explicitly asked to depict self-group overlap using a pictorial measure, self-ingroup overlap is consistently higher than self-outgroup overlap (Schubert & Otten, 2002), and some even conceptualize ingroup identification as the inclusion of the ingroup in the self (Tropp & Wright, 2001). Even constructs that are generally thought of as largely group-based, like social dominance orientation (Pratto et al., 1994), may in fact be primarily self-focused (Halali, Dorfman, Jun, & Halevy, 2017). Finally, ingroup helping can sometimes be driven by an expectation of reciprocity; even when the participant does not know the actions taken by others, they help more when they believe others' actions are also affecting them (Allidina et al., 2019; Gaertner & Insko, 2000; Rabbie et al., 1989; Rabbie & Lodewijkx, 1994). Overall, this research suggests that the line between self and ingroup is often a blurry one, lending support to the idea of helping the ingroup as a means to helping the self.

8.2. Limitations and future directions

Building on work demonstrating the existence of goal tradeoffs in intergroup contexts, here we made use of a novel decision-making task

to gain insight both on the kinds of tradeoffs that exist in these contexts and on the mechanisms and dynamics of these tradeoffs. The use of arbitrary minimal groups in this task allowed us to more fully control the intergroup environment, removing the possibility for existing group-based motivations to alter the processes at play. However, future work should examine the generalization of these processes to more externally valid social groups and explore the extent to which other motivational factors may play in. For example, existing social groups like those based on race or gender are likely to produce stronger group identification than our arbitrary minimal groups, which could potentially result in even stronger self-ingroup overlap and thus greater goal tradeoffs. Similarly, more hostile or competitive intergroup contexts may engender stronger motivations towards the outgroup, potentially producing greater self-outgroup tradeoffs than were seen here. Further, contexts in which an individual's identity is strongly "fused" with that of their group, leading them to make extreme sacrifices for the group (Swann et al., 2012; Swann et al., 2014), may shift the tradeoffs that are seen. Given that the self seems to be more motivationally primary than the ingroup even under high fusion (Heger et al., 2023), future research should examine how people in these more extreme circumstances choose to weigh and de-weigh their own outcomes for the sake of their group.

These studies examined multiple goal pursuit by including conditions where a primary self-related goal is "blocked" and cannot be pursued. A complementary approach would be to examine the pursuit of multiple goals when a primary goal has been fulfilled and thus no longer captures attention as strongly. For example, future research could manipulate opportunities for goal fulfillment early on to see whether those who have greater opportunity to help themselves are more likely to shift to helping their ingroup as trials progress. Such an approach could examine whether participants engage in a balancing act of multiple goals, choosing to upweight a previously secondary goal when they have already made substantial progress towards a primary motivation.

Finally, these studies were conducted on Canadian undergraduate students and online US-based participants, and thus suffer from the same issues around sample generalizability that plagues much of the field (Henrich, Heine, & Norenzayan, 2010). Future work will need to assess how the diversity of social environments and group norms outside of a North American context may alter or provide further nuance to these results. In addition, while some racial minority groups were well-represented in our samples, others were underrepresented. Racial diversity of samples will be especially important when examining how these cognitive processes generalize to real-world social groups, which may differ in their motivations and positions within group-based hierarchies. Thus, further considering the interplay of various motivations can add to our understanding of goal tradeoffs in intergroup contexts, with potentially important implications for reducing selfish or group-biased behaviours.

Open practices

Data and analysis code for all studies in this paper are available on the Open Science Framework at https://osf.io/ud7cr/?view_only=e42e1ce67b9340b08cd1a43cffe9352.

Authors' note

We thank Alice Wang, Victoria Chen, and Yuwen Chen for their help with data collection. We also thank the members of the Social Cognitive Science lab for their feedback on this work. This work was funded by the Social Sciences and Humanities Research Council (grant 506547) and the Natural Sciences and Engineering Research Council (grant RGPIN-2018-05946) to WAC. Data and analysis code for all studies in this paper are available on the Open Science Framework at https://osf.io/ud7cr/?view_only=e42e1ce67b9340b08cd1a43cffe9352.

ORCID iD authorship contribution statement

Suraiya Allidina: Conceptualization, Formal analysis, Investigation, Methodology, Visualization, Writing – original draft, Writing – review & editing. **William A. Cunningham:** Conceptualization, Funding acquisition, Methodology, Resources, Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jesp.2024.104625>.

References

- Aaldering, H., & Böhm, R. (2020). Parochial versus universal cooperation: Introducing a novel economic game of within- and between-group interaction. *Social Psychological and Personality Science*, 11(1), 36–45. <https://doi.org/10.1177/1948550619841627>
- Aaldering, H., Ten Velden, F. S., van Kleef, G. A., & De Dreu, C. K. W. (2018). Parochial cooperation in nested intergroup dilemmas is reduced when it harms out-groups. *Journal of Personality and Social Psychology*, 114(6), 909–923. <https://doi.org/10.1037/pspi0000125>
- Abbink, K., Brandts, J., Herrmann, B., & Orzen, H. (2010). Intergroup conflict and intra-group punishment in an experimental contest game conflict and intra-group punishment in an intergroup experimental contest game. *American Economic Review*, 100(1), 420–447. <https://doi.org/10.1257/aer.100.1.420>
- Abbink, K., Brandts, J., Herrmann, B., & Orzen, H. (2012). Parochial altruism in intergroup conflicts. *Economics Letters*, 117(1), 45–48. <https://doi.org/10.1016/j.econlet.2012.04.083>
- Ahmed, A. M. (2007). Group identity, social distance and intergroup bias. *Journal of Economic Psychology*, 28(3), 324–337. <https://doi.org/10.1016/j.joep.2007.01.007>
- Allidina, S., Arbuckle, N. L., & Cunningham, W. A. (2019). Considerations of mutual exchange in prosocial decision-making. *Frontiers in Psychology*, 10. <https://doi.org/10.3389/fpsyg.2019.01216>
- Balliet, D., Wu, J., & Dreu, C. K. W. D. (2014). *Ingroup favoritism in cooperation: A meta-analysis*. Psychological Bulletin. Psychological Bulletin Ingroup Favoritism in Cooperation, Advance online publication.
- Baron, J. (1997). The illusion of morality as self-interest: A reason to cooperate in social dilemmas. *Psychological Science*, 8(4), 330–335. <https://doi.org/10.1111/j.1467-9280.1997.tb00448.x>
- Baron, J. (2001). Confusion of group interest and self-interest in parochial cooperation on behalf of a group. *Journal of Conflict Resolution*, 45(3), 283–296. <https://doi.org/10.1177/0022002701045003002>
- Bates, D., Mächler, M., Bolker, B. M., & Walker, S. C. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1). <https://doi.org/10.18637/jss.v067.i01>
- Ben-Ner, A., Kramer, A., & Levy, O. (2008). Economic and hypothetical dictator game experiments: Incentive effects at the individual level. *The Journal of Socio-Economics*, 37(5), 1775–1784. <https://doi.org/10.1016/j.jsoec.2007.11.004>
- Böhm, R., Rusch, H., & Baron, J. (2018). The psychology of intergroup conflict: A review of theories and measures. *Journal of Economic Behavior and Organization*, 0, 1–16. <https://doi.org/10.1016/j.jebo.2018.01.020>
- Brueggemann, J., & Boswell, T. (1998). Realizing solidarity: Sources of interracial unionism during the great depression. *Work and Occupations*, 25(4), 436–482.
- Bürkner, P.-C. (2017). Brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80, 1–28. <https://doi.org/10.18637/jss.v080.i01>
- Camerer, C. F. (1989). An experimental test of several generalized utility theories. *Journal of Risk and Uncertainty*, 2(1), 61–104. <https://doi.org/10.1007/BF00055711>
- 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.
- De Cremer, D., & Van Vugt, M. (1999). Social identification effects in social dilemmas: A transformation of motives. *European Journal of Social Psychology*, 29(7), 871–893. [https://doi.org/10.1002/\(SICI\)1099-0992\(199911\)29:7<871::AID-EJSP962>3.0.CO;2-I](https://doi.org/10.1002/(SICI)1099-0992(199911)29:7<871::AID-EJSP962>3.0.CO;2-I)
- Dink, J., & Ferguson, B. (2018). EyetrackingR (R package version 0.1.8). <http://www.eyetracking-r.com>.
- de Dreu, C. K. W. (2010). Social value orientation moderates ingroup love but not outgroup hate in competitive intergroup conflict. *Group Processes & Intergroup Relations*, 13(6), 701–713. <https://doi.org/10.1177/1368430210377332>
- de Dreu, C. K. W., Dussel, D. B., & Ten Velden, F. S. (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(MAY), 1–9. <https://doi.org/10.3389/fpsyg.2015.00572>

- Fehr, E., Bernhard, H., & Rockenbach, B. (2008). Egalitarianism in young children. *Nature*, 454(7208), 1079–1083.
- Fernandez, J., & Kruglanski, A. (2018). The psychology of multiple goal pursuit: Choices, configurations, and commitments. *Journal of the Association for Consumer Research*, 4(1), 5–12.
- Fox, J., & Weisberg, S. (2019). *An R companion to Applied Regression* (3rd ed.). Sage Publications.
- Gaertner, L., & Insko, C. A. (2000). Intergroup discrimination in the minimal group paradigm: Categorization, reciprocation, or fear? *Journal of Personality and Social Psychology*, 79(1), 77–94.
- Gaertner, L., Sedikides, C., & Graetz, K. (1999). In search of self-definition: Motivational primacy of the individual self, motivational primacy of the collective self, or contextual primacy? *Journal of Personality and Social Psychology*, 76, 5–18.
- Gaertner, L., Sedikides, C., Luke, M., & Iuzzini, J. (2008). Hierarchy among selves: An implication for relations with persons versus groups. In H. A. Wayment, & J. J. Bauer (Eds.), *Transcending self-interest: Psychological explorations of the quiet ego* (pp. 127–135). American Psychological Association. <https://doi.org/10.1037/11771-012>.
- Gaertner, L., Sedikides, C., Luke, M., O'Mara, E. M., Iuzzini, J., Jackson, L. E., ... Quiping, W. (2012). A motivational hierarchy within: Primacy of the individual self, relational self, or collective self? *Journal of Experimental Social Psychology*, 48, 997–1013.
- Gaertner, L., Sedikides, C., Vevea, J., & Iuzzini, J. (2002). The “I,” the “we,” and the “when”: A meta-analysis of motivational primacy in self-definition. *Journal of Personality and Social Psychology*, 83, 574–591.
- Gillis, M. T., & Hettler, P. L. (2007). Hypothetical and real incentives in the ultimatum game and Andreoni's public goods game: An experimental study. *Eastern Economic Journal*, 33(4), 491–510. <https://doi.org/10.1057/eej.2007.37>
- Halali, E., Dorfman, A., Jun, S., & Halevy, N. (2017). More for us or more for me? Social Dominance as Parochial Egoism. *Social Psychological and Personality Science*. <https://doi.org/10.1177/1948550617732819>
- 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>
- Halevy, N., Weisel, O., & Bornstein, G. (2011). “In-group love” and “out-group hate”. In , Vol. 2. *Repeated interaction between groups*. <https://doi.org/10.1002/bdm>
- Heger, A., Voorhees, V. A., Porter, B., & Gaertner, L. (2023). Does identity fusion moderate the motivational primacy of the individual self? *Social Psychological and Personality Science*, 14(8), 934–941. <https://doi.org/10.1177/19485506221141144>
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2–3), 61–83. <https://doi.org/10.1017/S0140525X0999152X>
- Kramer, R. M., & Brewer, M. B. (1984). Effects of group identity on resource use in a simulated commons dilemma. *Journal of Personality and Social Psychology*, 46(5), 1044–1057. <https://doi.org/10.1037/0022-3514.46.5.1044>
- Kruglanski, A. W., Chernikova, M., Babush, M., Dugas, M., & Schumpe, B. M. (2015). The architecture of goal systems: Multifinality, equifinality, and counterfinality in means-end relations. In , Vol. 2. *Advances in motivation science* (pp. 69–98). <https://doi.org/10.1016/bs.adms.2015.04.001>
- Kruglanski, A. W., Köpetz, C., Bélanger, J. J., Chun, W. Y., Orehek, E., & Fishbach, A. (2013). Features of multifinality. *Personality and Social Psychology Review*, 17(1), 22–39. <https://doi.org/10.1177/1088868312453087>
- Kruglanski, A. W., Shah, J. Y., Fishbach, A., Friedman, R., Chun, W. Y., & Sleeth-Keppler, D. (2018). A theory of goal systems. In *The motivated mind*. <https://doi.org/10.4324/9781315175867-6>
- Kühberger, A., Schulte-Mecklenbeck, M., & Perner, J. (2002). Framing decisions: Hypothetical and real. *Organizational Behavior and Human Decision Processes*, 89(2), 1162–1175. [https://doi.org/10.1016/S0749-5978\(02\)00021-3](https://doi.org/10.1016/S0749-5978(02)00021-3)
- Leach, C. W., Spears, R., Branscombe, N. R., & Doosje, B. (2003). Malicious pleasure: Schadenfreude at the suffering of another group. *Journal of Personality and Social Psychology*, 84(5), 932.
- Leach, C. W., Van Zomeren, M., Zebel, S., Vliek, M. L. W., Pennekamp, S. F., Doosje, B., ... 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>
- Lenth, R. V. (2021). Emmeans: Estimated marginal means, aka least-squares means. R package version 1.7.0 <https://CRAN.R-project.org/package=emmeans>.
- Mahoney, M. R. (1999). Constructing solidarity: Interest and White workers. *University of Pennsylvania Journal of Labor and Employment Law*, 2, 747–771.
- Murphy, R. O., & Ackermann, K. A. (2013). Social value orientation: Theoretical and measurement issues in the study of social preferences. *Personality and Social Psychology Review*, 18(1), 1–29. <https://doi.org/10.1177/1088868313501745>
- Orehek, E., & Vazeou-Nieuwenhuis, A. (2013). Sequential and concurrent strategies of multiple goal pursuit. *Review of General Psychology*, 17(3), 339–349. <https://doi.org/10.1037/a0032584>
- Pratto, F., Sidanius, J., Stallworth, L. M., & Malle, B. F. (1994). Social dominance orientation: A personality variable predicting social and political attitudes. *Journal of Personality and Social Psychology*, 67(4), 741–763. <https://doi.org/10.1037/0022-3514.67.4.741>
- Purzycki, B. G., & Lang, M. (2019). Identity fusion, outgroup relations, and sacrifice: A cross-cultural test. *Cognition*, 186(November), 1–6. <https://doi.org/10.1016/j.cognition.2019.01.015>
- R Core Team. (2021). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. <https://www.R-project.org/>.
- Rabbie, J. M., & Lodewijkx, H. F. M. (1994). Conflict and aggression: An individual-group continuum. *Advances in Group Processes*, 11, 139–174.
- Rabbie, J. M., Schot, J. C., & Visser, L. (1989). Social identity theory: A conceptual and empirical critique from the perspective of a behavioural interaction model. *European Journal of Social Psychology*, 19(3), 171–202. <https://doi.org/10.1002/ejsp.2420190302>
- Rahal, R. M., Fiedler, S., & De Dreu, C. K. W. (2020). Prosocial preferences condition decision effort and Ingroup biased generosity in intergroup decision-making. *Scientific Reports*, 10(1), 1–11. <https://doi.org/10.1038/s41598-020-64592-2>
- van Rij, J., Wieling, M., Baayen, R., & van Rijn, H. (2020). *itsadug: Interpreting time series and autocorrelated data using GAMMs (R package version 2.4)*.
- Schubert, T. W., & Otten, S. (2002). Overlap of self, Ingroup, and outgroup: Pictorial measures of self-categorization. *Self and Identity*, 1(4), 353–376. <https://doi.org/10.1080/152988602760328012>
- Shah, J. Y., Friedman, R., & Kruglanski, A. W. (2002). Forgetting all else: On the antecedents and consequences of goal shielding. *Journal of Personality and Social Psychology*, 83(6), 1261–1280. <https://doi.org/10.1037/0022-3514.83.6.1261>
- Smith, R. H., Powell, C. A., Combs, D. J., & Schurtz, D. R. (2009). Exploring the when and why of schadenfreude. *Social and Personality Psychology Compass*, 3(4), 530–546.
- 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>
- Swann, W. B., Buhrmester, M. D., Gómez, A., Jetten, J., Bastian, B., Vázquez, A., ... Zhang, A. (2014). What makes a group worth dying for? Identity fusion fosters perception of familial ties, promoting self-sacrifice. *Journal of Personality and Social Psychology*, 106(6), 912–926. <https://doi.org/10.1037/a0036089>
- Swann, W. B., Gómez, Á., Seyle, D. C., Morales, J. F., & Huici, C. (2009). Identity fusion: The interplay of personal and social identities in extreme group behavior. *Journal of Personality and Social Psychology*, 96(5), 995–1011. <https://doi.org/10.1037/a0013668>
- Swann, W. B., Jetten, J., Gómez, Á., Whitehouse, H., & Bastian, B. (2012). When group membership gets personal: A theory of identity fusion. *Psychological Review*, 119(3), 441–456. <https://doi.org/10.1037/a0028589>
- Tajfel, H., Billig, M. G., & Bundy, R. P. (1971). Social categorization and intergroup behaviour. *European Journal of Social Psychology*, 1(2), 149–178.
- Thielmann, I., & Böhm, R. (2016). Who does (not) participate in intergroup conflict? *Social Psychological and Personality Science*, 7(8), 778–787. <https://doi.org/10.1177/1948550616660160>
- Tropp, L. R., & Wright, S. C. (2001). Ingroup identification as the inclusion of ingroup in the self. *Personality and Social Psychology Bulletin*, 27(5), 585–600. <https://doi.org/10.1177/0146167201275007>
- White, I. K., Laird, C. N., & Allen, T. D. (2014). Selling out?: The politics of navigating conflicts between racial group interest and self-interest. *American Political Science Review*, 108(4), 783–800. <https://doi.org/10.1017/S000305541400046X>
- Wieling, M. (2018). Analyzing dynamic phonetic data using generalized additive mixed modeling: A tutorial focusing on articulatory differences between L1 and L2 speakers of English. *Journal of Phonetics*, 70, 86–116. <https://doi.org/10.1016/j.wocn.2018.03.002>
- Wood, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, 73(1), 3–36. <https://doi.org/10.1111/j.1467-9868.2010.00749.x>
- Yamagishi, T., & Mifune, N. (2016). Parochial altruism: Does it explain modern human group psychology? *Current Opinion in Psychology*, 7, 39–43. <https://doi.org/10.1016/j.copsyc.2015.07.015>