Efficiency is Prioritised Over Fairness When Distributing Joint Actions

James W.A. Strachan[1,2] & Georgina Török[1]

1. Central European University, Budapest, Hungary
2. Italian Institute of Technology, Genova, Italy

Corresponding author:

Name: James W.A. Strachan

Email: james.wa.strachan@gmail.com

Address:

Fondazione Istituto Italiano di Tecnologia,

Via Enrico Melen 83,

16152 Genova

Italy

July 2020: *Manuscript as accepted for special issue in* Acta Psychologica *on Current Issues in Joint Action Research Edited by Robrecht van der Wel, Cristina Becchio, Arianna Curioni & Thomas Wolf.*

People typically move rationally, aiming to reach their goals in the most efficient way possible, given their beliefs and the environment's constraints. This is a key assumption of action observation systems tasked with understanding others' actions, which adopt a teleological stance based on this rationality principle (Gergely & Csibra, 2003). Deviations from optimality are presumed to serve a rational function, such as communication (Csibra et al., 2003; Dockendorff et al., 2019; Pezzulo et al., 2013; Royka et al., 2018), and action rationality is necessary for the attribution of goals (Southgate et al., 2014).

This assumption is borne out by evidence that people perform individual actions that conform to the rationality principle: their movement planning often minimises factors such as movement time or effort (Lyons et al., 2006; Rosenbaum et al., 2001), or distance covered by the hand (Török et al., 2019), and favours the recruitment of limb segments that best fit a task's demands (Rosenbaum et al., 1991). Furthermore, people will initially perform costly or uncomfortable actions, such as adopting grasps in awkward joint angles, in order to maximise the end-state comfort of their action (Cohen & Rosenbaum, 2004; Meyer et al., 2013; Rosenbaum et al., 1990). Such findings suggest that people plan their actions prospectively to minimise total expected action costs, that is, to maximise their efficiency.

When pursuing a shared goal with someone, the action affordances, expected costs, and constraints of the partner must be incorporated into a person's action plan, if efficiency of the joint action were to be maximised. In support of this idea, actors were found to co-represent the task demands and environmental constraints of their partner during coordinated movement (Schmitz et al., 2017; Vesper et al., 2013), and plan sequential grasping and passing actions to minimise discomfort for their partner, e.g. by avoiding uncomfortable postures or difficult movements for their partner (Meyer et al., 2013; Ray et al., 2017; Ray & Welsh, 2011). To complement these results, participants in economic games made payoff choices that maximised the total expected utility of a virtual dyad they belonged to (Colman et al., 2008). Crucially, individuals in joint actions are willing to incur additional motor costs (perform longer, more difficult actions) when these lead to a lower total cost over the whole action sequence (Török et al., 2019; see also the shared-effort model proposed by Santamaria & Rosenbaum, 2011), suggesting that people prioritise joint or *coefficiency* over their own individual efficiency during joint action.

In Török and colleagues' (2019) study, participants had to pass an object to a co-actor on a touchscreen by dragging it with their index fingers. In each trial, moving the object entailed completing a joint movement sequence divided into two halves in the middle of the screen, where co-actors transferred the object to each other in one of two potential transfer points. On a trial-by-trial basis, participants decided between the two transfer points, i.e., the participant initiating the sequence chose an overall action plan for the dyad. Movement costs were operationalised as proportional to the distance that the hand covered on the screen, and they were manipulated by adding walls of various sizes to the screen layout, which blocked the movement of the object. Participants could either choose to minimise their own individual action costs by picking the transfer point closer to them, or to minimise their partner's individual action costs by choosing the transfer point farther from themselves (and closer to their partner). Depending on the manipulation of the expected movement costs in a given trial, minimising the overall movement costs of the dyad coincided either with minimising the decision-making participant's own individual costs or their partner's individual costs. The results of three joint experiments indicated that people chose action plans that minimised the overall costs of an action sequence, when possible. That is, they maximised coefficiency - regardless of whether such a decision also maximised their own or their partner’s individual efficiency, even at a cost to themselves in the latter case.

Compromising one’s own individual efficiency for the benefit of the group can be risky, however. Imagine two friends who live a considerable walking distance from each other agree to meet in the city: they have two preferred cafés where they can meet, one which is far away from both of their homes, and the other which is close to one friend’s house but far from the other’s. In the latter case the sum walking distance of the friends to the café is shorter than in the former case. To minimise the total walking distance, the optimal decision would be to go to the café closest to one friend, but this means that one friend is investing much greater effort in the meeting than the other – this could become particularly salient over repeated interactions where they always meet in the same place. In order to counter any feelings of resentment or unfairness, the friends may decide to visit the more distant meeting point in order to ensure an equitable investment of effort.

People are very sensitive to fairness. Infants expect agents to distribute resources equally (Schmidt & Sommerville, 2011; Sloane et al., 2012) and, in cases of asymmetric task contributions, expect an equitable allocation of rewards (such that they expect rewards to be appropriate relative to the amount of work invested; Wang & Henderson, 2018). Participants in economic games reject unfair reward distributions, even when they benefit from the unfairness, or rejection comes at an individual cost (Blake & McAuliffe, 2011; Dawes et al., 2007). Furthermore, there is evidence that fairness may be intrinsically rewarding (Tabibnia & Lieberman, 2007).

However, many studies of fairness look at instances where participants themselves are the recipients or distributors of the outcomes of fair or unfair behaviour – either when interacting with a free-riding participant (Fehr & Gächter, 2002; Lepine & Van Dyne, 2001; Melis et al., 2013; Taggar & Neubert, 2008) or when being offered an unfair portion of the rewards (Brethel-Haurwitz et al., 2016; Civai et al., 2010; Clark et al., 2017; Gu et al., 2016) – and measures focus primarily on the distribution of rewards or resources (Dawes et al., 2007). It remains an open question whether participants have a similar drive for fairness in motor task distribution, and whether their decisions to distribute a task reflect this. Importantly, it is also unclear how fairness is treated rationally in task allocation. While the drive to maximise coefficiency may play a large role in how people distribute tasks in joint actions due to its pragmatic cost-saving effects, it remains to be seen how much importance is placed on being fair to their co-actor when distributing these tasks. If people do consider fairness when allocating tasks in joint actions, this suggests that they incorporate the potential social consequences of being perceived as fair or unfair into rational utility estimates of planned joint actions.

The current study explored whether participants would continue to prioritise coefficiency in cases where being maximally efficient in a joint action would require asymmetric contributions from the two co-actors, even when a less coefficient but fairer means of completing the joint action was also available to the dyad. In two experiments, participants performed the first part of a joint passing action with a virtual partner (moving an object to a target by passing it to a partner who then – they believed – completed the movement). In both experiments, participants had a choice between a symmetrical route and an asymmetrical route that could be overall longer, shorter, or equal in length to the symmetrical route. In the symmetrical case, each actor would move the object the same distance, whereas in the asymmetrical case, one actor would invest more effort in completing the joint action than their co-actor. If participants were driven to maximise coefficiency, they should consistently pick the overall shortest route, regardless of its fairness. On the other hand, if they were driven to prioritise fairness, they would be more likely to consistently pick the symmetrical route.

In Experiment 1, on a given trial the asymmetrical route could demand greater investment from either the participant (own path longer condition) or the partner (other path longer condition). Such counterbalancing means that participants could achieve fairness by selecting a coefficiency-maximising strategy because choosing the shortest overall route throughout the experiment would lead to them completing as many long distances as their partner. To control for this potential explanation, in Experiment 2 we made the asymmetrical path always shorter for the participant than the symmetrical path (individually efficient), so that the only way to choose a coefficiency-maximising strategy would be to force the partner to incur relatively more costs overall.

Materials and Methods

Participants

In Experiment 1, 50 participants were recruited. Due to exclusion criteria (see below), 26 of these were excluded from the final sample, leaving 24 in total (14 female, 10 male; M_{age}=29.29yrs). In Experiment 2, 36 participants were recruited with 12 exclusions, leaving 24 in total (8 female, 16 male; M_{age}=25.25yrs). Participants were recruited from the Central European University SONA recruitment service. Participants received 1500 HUF (approx. €4.60) worth of vouchers that could be redeemed at a variety of local shops. All experiments in this study were approved by the United Ethical Review Committee for Research in Psychology (EPKEB). Participants were recruited in pairs and run concurrently, under the impression that they were playing with their partner in different rooms. Participant pairs were questioned when they first arrived and only participants pairs who reported that they did not know each other took part in the study.

Stimuli and Design

Experiment 1

Origin and target positions (where the object started from and moved towards, respectively) remained the same across all trials, positioned vertically 30° apart in the centre of the screen. The origin and target were connected by two routes forming an irregular quadrilateral, and each route consisted of two individual paths separated by a passing junction. One side (left/right counterbalanced across trials) was the symmetrical side: the two individual movement paths created an isosceles triangle with the central vertex between the origin and target (see Figure 1A; examples in Figure 1B). The total length of this symmetrical side could be one of nine total lengths, three of which were designated short (39°, 41.25°, and 43.5°, or 41.25°±2.25), medium (48.125°±2.625), or long (55°±3). The individual movements required from each participant were identical. The asymmetrical side was calculated in relation to the symmetrical side, and could be either longer (symmetrical-coefficient, in that the symmetrical side was also the more efficient path overall), shorter (asymmetrical-coefficient, in that prioritising efficiency and prioritising fairness would lead to different outcomes), or equal in length to the fair side (the two sides are equally coefficient and so coefficiency is not a consideration). In total, 72 possible configurations were used (12 configurations in each condition: symmetrical-coefficient, equal, and asymmetrical-coefficient; with participants' own or the other path longer).

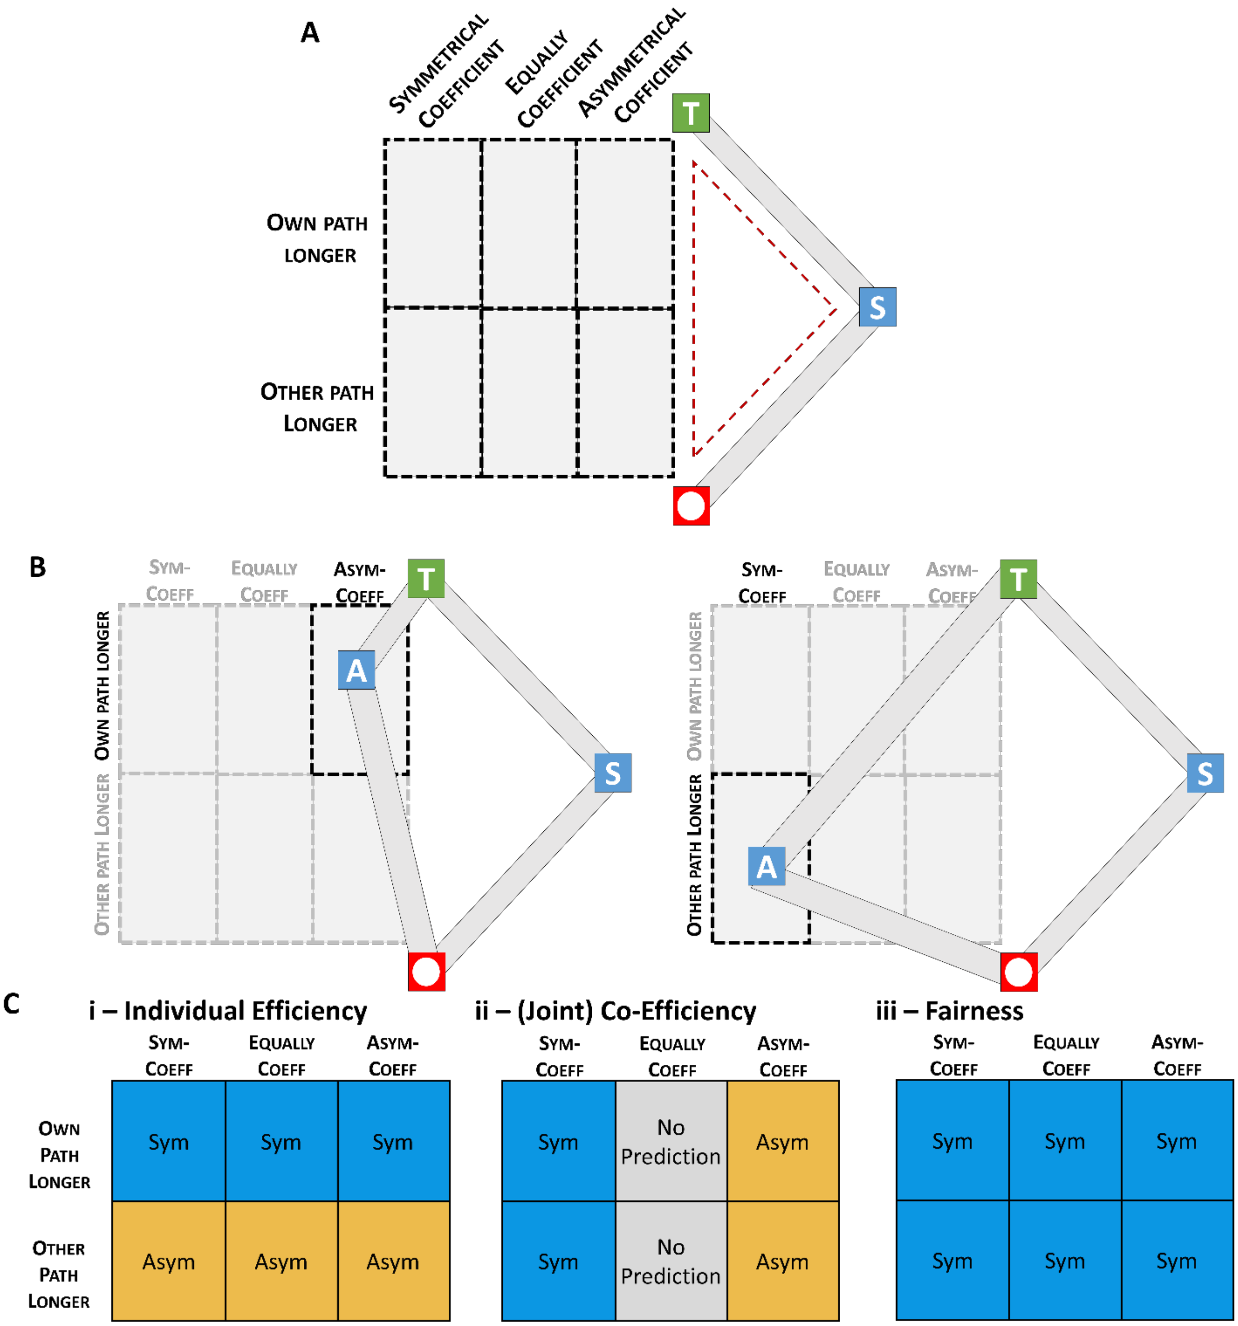

Figure 1. *Design and predictions of Experiment 1. In all figures, labels on the passing junctions (i.e. S and A) and on the target (T) and origin (O) are shown for illustrative purposes and were not presented to participants. A. All stimuli were generated with a symmetrical side (S, requiring equal input from both parties) and an asymmetrical side (A) that could fall into one of six conditions that were either symmetrical-coefficient, equally coefficient, or asymmetrical-coefficient, and the asymmetrical side could be biased such that the participant had the longer movement (Own path longer) or their partner had the longer movement (Other path longer). Note that in Experiment 2 there were no Own path longer trials, and so the A junction would only appear in the lower half of the grid. B. Two example configurations. Condition boxes and labels are shown here for illustrative purposes and were not visible to participants. C. Predictions of three potential decision-making strategies as to which path participants would select:*

(i) an Individual Efficiency-maximising strategy would result in participants seeking to minimise their own motor costs and so pick the shortest route for themselves; (ii) a Coefficiency-maximising strategy would result in participants choosing the shortest route overall, regardless of whether this incurs greater costs for themselves; (iii) a Fairness maximising strategy would result in participants always choosing the fair path.

Each configuration was presented with the fair side on the left and on the right, and each configuration and orientation appeared twice throughout the experiment, resulting in 288 trials per participant. All stimuli and instructions were presented on a black background.

Experiment 2

Experiment 2 was identical to Experiment 1, except that there were no trials where the asymmetrical side put the participant at a disadvantage (only the Other-path longer conditions, see Figure 1A,D). We implemented this modification because with both conditions, participants could rationalise that a coefficiency-maximising strategy was ultimately fair, since it would put themselves at a disadvantage as often throughout the whole task as the other (see Figure 1D.ii). As such, if participants chose to maximise coefficiency in Experiment 2 they would be consistently forcing their partner to incur greater motor costs than themselves.

Due to using half the number of configurations in this experiment, we doubled the number of repetitions (from two to four) to retain the same number of trials as in Experiment 1.

Procedure

The experiment consisted of three stages. In the first stage, participants completed the main experiment passing objects. Once that ended, participants were taken to a questionnaire to answer a series of questions about their experience. Finally, participants were asked to complete a perceptual judgement task.

Main experiment. Participants were recruited in pairs and tested in separate testing rooms. They were told that they were playing with the participant in the other room and had been randomly selected to be the leader in the interaction - in fact they were playing with a virtual partner and both participants were the leader. This meant that they would complete the first part of the action and their partner would complete the second part. Participants were shown an example configuration with both sides symmetrical and equal and told that their task was to pass the item (a white circle) from the red origin to the green target site by passing it to their partner at one of the two blue junctions.

The structure of a single trial is shown in Figure 2. The mouse cursor was visible from the beginning of the trial and appeared centred over the item at the origin at the same time. Participants clicked on the item and dragged it along one of the grey lines to one of the passing junctions. The item would only move if it overlapped with the grey line or one of the squares, preventing any possible deviations from the assigned paths. Once the participant reached one of the passing junctions, the item would lock in place and the mouse cursor would disappear. After a brief delay the item would turn yellow, which participants were told indicated that the ‘partner’ had clicked on it. The item remained yellow and did not move until the ‘partner’s’ action was finished, at which point it appeared at the target site for 1,000ms before a prompt to advance to the next trial by pressing SPACE was displayed.

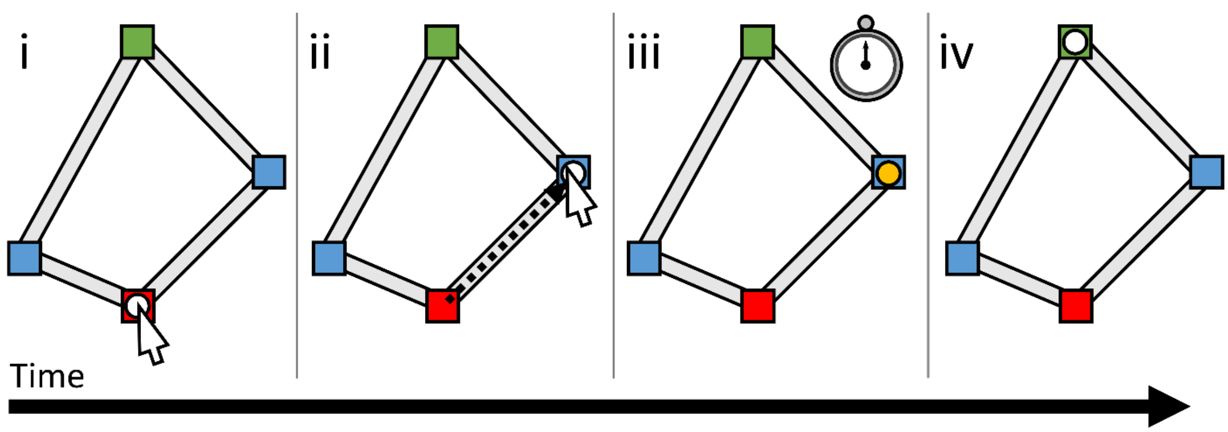

Figure 2. *Example trial sequence: (i) Trial starts with mouse cursor position over the item at origin point. (ii) Participant must click on the item and drag it to one of the passing junctions – either the symmetrical or asymmetrical side. Dotted arrow shows dragging motion.*

The delays before the item turned yellow and before it appeared in the target position were calculated as a function of the average of participants' own performance on the preceding three trials and the distance that the 'partner' would have to move. As such, the partner was no better or worse at the task than the participant and appeared to move realistically. In Experiment 1, the time between the item turning yellow and appearing in the target position was calculated using the following formula: $\mu(x_{[t-3,t]})/h * r$, where x represents the participant's movement duration from picking up the item to depositing it on trial t, h is a constant representing the total vertical distance from origin to target (30°), and r indicates the distance of the partner's own path that the item must travel. Given the high number of participants in Experiment 1 who demonstrated suspicion of the virtual partner being the human in the other room, and that many participants reported the partner's actions as being noticeably fast, in Experiment 2 we slowed down the partner's movements by changing h to $h/2$.

Questionnaire. On the final screen of the main experiment, participants were given a link to a Google Form questionnaire. In this questionnaire we asked whether participants believed at any point that their partner was not a real person. Participants were excluded if they responded Yes to this question (11 in Experiment 1, plus 3 due to a technical error resulting in missing questionnaire data; 7 in Experiment 2).

Perceptual judgements. To ensure that results were not driven by participants being unsure as to which path was more or less efficient, the final section of the experiment involved participants making decisions about the configurations they had seen during the main experiment. Participants were asked to judge whether the two sides of the shape were equal (press SPACE) or if one was longer (left or right arrow keys). We removed any participants who made >30% errors on judging either congruent or incongruent trials (12 in Experiment 1, although see footnote 1 in Results section; 5 in Experiment 2).

Predictions

The predictions for each potential decision-making strategy in both experiments are illustrated in Figure 1C. We predicted that if participants followed an Individual Efficiency-maximising strategy, they would opt for the shortest path for themselves, regardless of its length relative to the partner's path length (i.e. its 'fairness'), or to the total path length (i.e. its coefficiency). Were participants to follow a Coefficiency-maximising strategy, aiming to minimise the overall path lengths to be covered by themselves and their partner, we expected participants to choose the symmetrical path in the symmetrical-coefficient trials, and asymmetrical in the asymmetrical-coefficient trials. Finally, we predicted that if participants followed a Fairness-maximising strategy, then we would observe a majority of symmetrical path choices, regardless of their (co-)efficiency or the identity of the actor that would be disadvantaged by an unfair task distribution in a given trial.

Data analysis

The decisions of whether to take the symmetrical or asymmetrical path were analysed using a logistic regression. In Experiment 1, this analysis included trial condition (equally coefficient as baseline vs. symmetrical-/asymmetrical-coefficient) and relative individual path length (own path longer vs. other path longer) as independent variables and decision outcome (symmetrical vs. asymmetrical) as the dependent variable. In Experiment 2, we only included trial condition as the independent variable.

Data were analysed in R, using the *glm* function.

Results

Experiment 1

The results of participants' decisions of which path to take in Experiment 1 are shown in the top row of Figure 3. On symmetrical-coefficient trials, as expected by both the coefficiency- and fairness-maximising strategies, participants chose the symmetrical option most frequently. On asymmetrical-coefficient trials, however, people chose the asymmetrical route more frequently than the symmetrical route, indicating a coefficiency-maximising strategy.

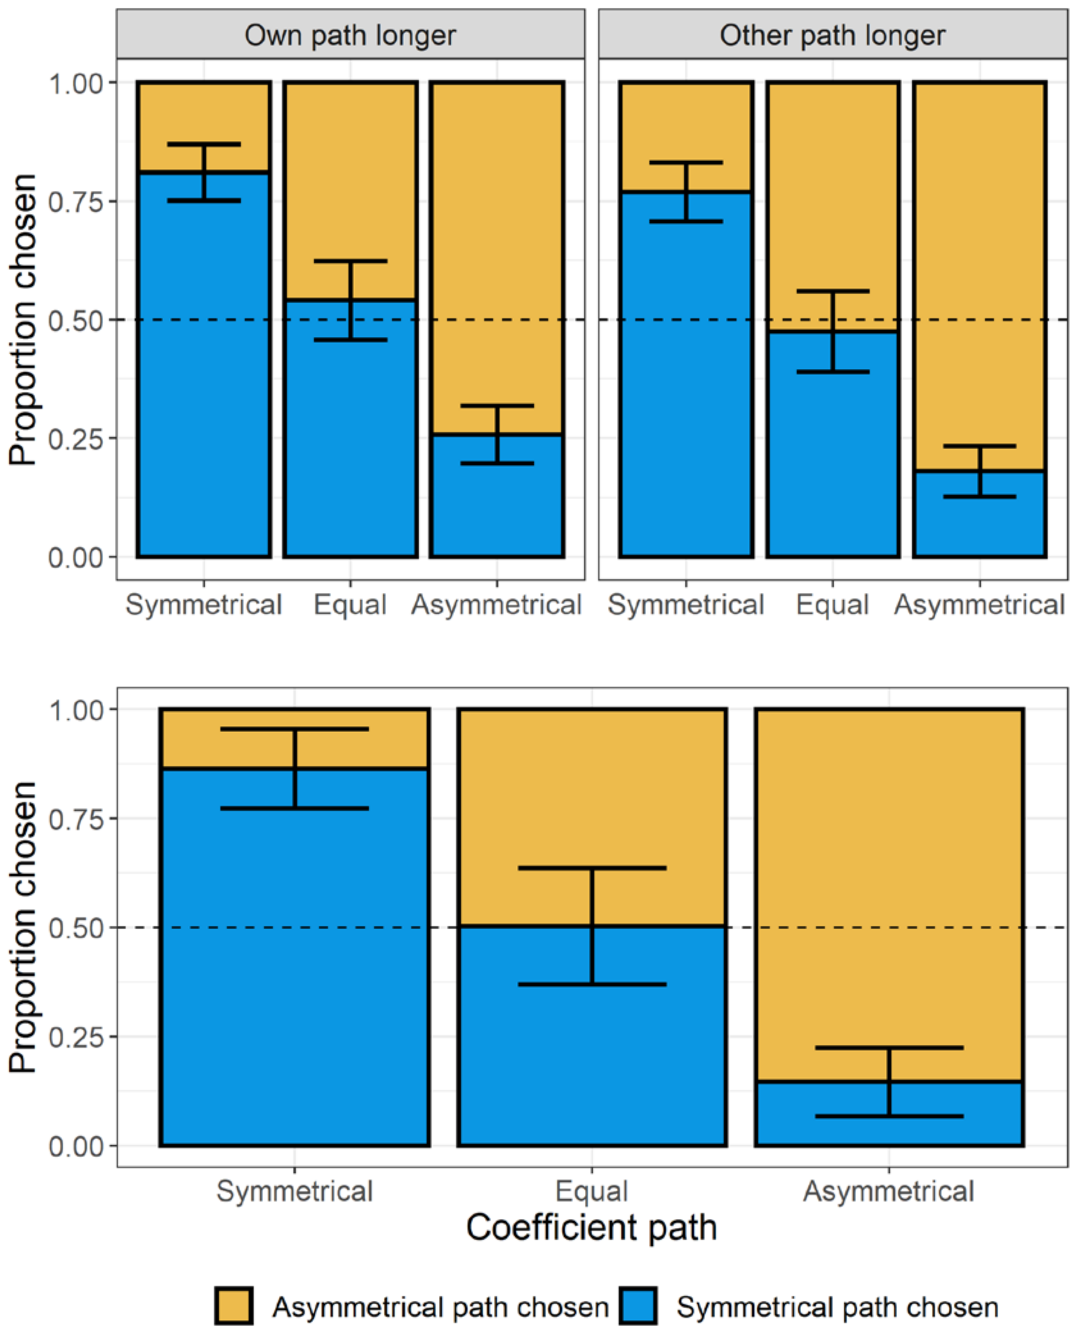

Figure 3. *Results of Experiments 1 (top row) and 2 (bottom row) as proportion of trials on which participants chose the asymmetrical (light/orange; top bars) and symmetrical (dark/blue; bottom bars) routes in each condition. Error bars show 95% confidence intervals.*

The results of the logistic regression found that participants were significantly more likely to pick the symmetrical route over the asymmetrical route when the symmetrical route was the coefficient option, compared to when the two paths were equally coefficient (β=-1.30, SE=0.09, z=-14.25, p<.001). Furthermore, participants were significantly more likely to pick the asymmetrical route over the symmetrical route when the asymmetrical route was the coefficient option than when the two paths were equally coefficient (β=1.41, SE=0.10, z=14.60, p<.001).

There was also a slight bias towards maximising individual efficiency. On trials where the two routes were equally coefficient, participants were significantly more likely to pick the symmetrical route if their own path was shorter on this than on the asymmetrical route (Figure 3, top left, middle stack, dark/blue bar extends slightly beyond 50% chance level; β=-0.26, SE=0.08, z=-3.17, p=.002). Participants were also non-significantly more likely to pick the asymmetrical route if their own path on this route was shorter than on the symmetrical route, all else being equal (Figure 3, top right, middle stack, light/orange bar extends slightly beyond 50% chance level; β=0.10, SE=0.06, z=1.71, p=.088). This indicates that, when the two routes were equally coefficient, participants were slightly more likely to choose the shorter path for themselves.

However, when the two routes differed in terms of coefficiency, there was no evidence of any interaction between whether the symmetrical or asymmetrical route was longer overall and whose path was longer within the asymmetrical route. When the symmetrical path was more coefficient than the asymmetrical path, participants were as likely to choose the symmetrical path when it was longer for themselves as when it was shorter (Figure 3; top left vs. top right; left stacks; β=0.02, SE=0.13, z=0.14, p=.890). When the asymmetrical path was more coefficient than the symmetrical, participants were as likely to choose the asymmetrical path when it was longer for themselves as when it was shorter (Figure 3; top left vs. top right; right stacks; β=-0.19, SE=0.13, z=-1.45, p=.148). As such, while participants would prioritise their own individual efficiency when they could, they were willing to incur additional movement costs if this maximised the efficiency of the joint action[1]. On the other hand, when the asymmetrical route was shorter than the symmetrical route, the symmetry of the longer route did not interfere with this coefficiency-maximising strategy, suggesting that participants were not willing to compromise overall efficiency to ensure a fair contribution from both actors.

Experiment 2

In Experiment 2 we explored whether the results of Experiment 1 could be driven by the inclusion of trials where participants' own path was longer. In this experiment, participants could only maximise coefficiency by being consistently unfair to their partner. Despite this, the results were strikingly similar to Experiment 1 (see the bottom row of Figure 3). In this experiment, on trials where the two routes were equally coefficient, the participant's path on the symmetrical route was always longer than on the asymmetrical route. If participants in Experiment 2 showed an individual efficiency-maximising bias on equally coefficient trials (as observed in Experiment 1), then they would consistently choose the asymmetrical route over the symmetrical on equally coefficient trials (Figure 3, bottom row, middle stack). However, in this experiment there was no evidence of such an individual efficiency-maximising strategy, as participants' choices did not differ from 50% chance (β=0.01, SE=0.04, z=0.25, p=.803).

However, participants were once again significantly more likely to choose the symmetrical route over the asymmetrical route when the symmetrical route was more coefficient (β=1.83, SE=0.07, z=24.91, p<.001), and they were more likely to choose the asymmetrical route over the symmetrical route when the asymmetrical route was more coefficient (β=-1.79, SE=0.07, z=-24.66, p<.001). This pattern of results – choosing the shortest overall path even though this resulted in the partner investing consistently more effort across the experiment –

indicates that participants used a coefficiency-maximising strategy instead of a fairness-maximising strategy.

Suspicious participants

Participants in Experiments 1 and 2 were asked whether they held any suspicion during the experiment that their partner was not a real partner. Participants were excluded if they reported yes, since if participants had been inclined to be fair to their partners we expected that this effect would have been specific to interacting with a human partner. However, given that we have seen overwhelming evidence for a coefficiency-maximising strategy, we decided to look closer at these excluded participants to see if this reflects a general preference for overall efficiency, or if participants who believe they are interacting with a computer would be more likely to use an individual efficiency-maximising strategy.

We looked at participants excluded from Experiment 1, who demonstrated suspicion of their partner but who satisfied the perceptual judgement criterion (scoring above 70% or, given that some participants misunderstood the instructions, below 30%) and found that of 11 excluded participants, 10 were suitable for closer scrutiny. Although we do not report statistical analysis of such a small, post-hoc sample, participants showed a pattern of data remarkably consistent with those who were included: when the symmetrical path was more efficient overall, participants chose it over the asymmetrical path when it was shorter for themselves (86.46% of all trials across participants), and also when it was the longer path for themselves (85.83% of all trials). When the asymmetrical path was more efficient overall, participants chose it more when it was shorter for themselves (80.83% of all trials), and also when this was the longer path (81.67% of all trials). On trials where both sides were equally coefficient, participants showed little preference for or against the symmetrical path (they chose the symmetrical path on 56.67% of all trials when the path was shorter for themselves, and on 52.50% of all trials when it was longer for themselves).

In Experiment 2, only four participants who demonstrated suspicion of their partner satisfied the perceptual judgement criterion, but even then these participants chose to maximise coefficiency, both when this meant choosing the shorter path for themselves (asymmetrical: 78.13%) and choosing the longer path for themselves (symmetrical: 93.75%). While we must avoid overinterpreting such limited data, the decisions made by suspicious participants in Experiments 1 and 2 may suggest that participants tend to maximise the overall efficiency at the expense of their individual efficiency, even if they do not necessarily believe that their partner is an intentional agent.

We find no evidence in this exploration of suspicious participants' data that they use substantially different strategies from participants who believe they are interacting with a real person. This could indicate that the drive to maximise coefficiency reflects a general-purpose strategy that is not bound to an interaction with an intentional agent. While this study was not designed to address this question of how beliefs about a partner's agency affect strategy choice (not least because we cannot be sure when participants' suspicions were first raised during the experiment), this finding poses some interesting questions for future research into how beliefs about a partner such as attributions of agency may affect rational decision making.

Discussion

Across two experiments, we consistently found that participants chose the shortest overall route when passing items to a partner, even when this resulted in unfair task distributions. In Experiment 1, we replicated previous results showing that participants were willing to incur individual movement costs to maximise the coefficiency of a dyad (Török et al., 2019). In

Experiment 2, where maximising coefficiency necessarily forced their partner to invest more effort than themselves, participants continued to choose to maximise coefficiency, indicating that fairness did not affect decision outcomes during rational decision-making of task distribution in a joint action.

These findings suggest that participants plan their actions considering the instrumental, physical costs – in this case movement costs associated with precise action that vary according to distance – and they take into account the total expected cost experienced by the pair. However, the within-dyad asymmetry between these costs, i.e. the unfairness of the task distribution, does not appear to affect participants' decisions. This suggests that participants are equally willing to sacrifice a partner's individual efficiency as they are their own, in order to maximise the total efficiency of the joint action.

It is possible that fairness was not considered in participants' decision because there was no chance for reciprocity in our design: participants were assigned to a 'leader' role that meant they made all decisions for a virtual partner who never had the opportunity to reciprocate. In fact, under such conditions the 'leader' may feel a pressure or responsibility to maximise the group's coefficiency that overrides concerns of unfairness. However, convergent evidence from other experiments that did include an opportunity for reciprocity of decision-making with a real, live partner support the finding that participants prioritise coefficiency over fairness (Török et al., submitted). The study by Török et al. (submitted) primarily focused on the computation of joint action costs and only considered a fairness-based decision-making process as an alternative hypothesis to coefficiency. Here, we designed our task specifically to address this question, and successfully replicated their pattern of results.

One may argue that the costs associated with the movements required in this study were too small for participants to be sensitive to their relative asymmetries. However, the consistent selection of coefficient paths indicates that these costs are indeed computed, and points once more to coefficiency being privileged over fairness. It is possible that fairness may affect decision-making more when there are substantial action costs or asymmetries, or when the unfairness of the task distribution is made more salient than the efficiency. Capraro and Rand (2018) found that, in economic games that pit equity against efficiency of outcomes, participants' decisions can be reliably swayed by whichever is framed as the morally right decision. Although the current study did not involve any explicit moral framing, we think that this question would be worth exploring in future research.

Investigating the role of fairness in task distribution, in particular how the social consequences of unfairness are balanced within rational models of joint action, offers a rich avenue for future research. In particular, as well as investigating the factors that might lead to the prioritisation of fairness over efficiency, it is important to examine the similarities and differences across fairness considerations in task distribution and resource distribution. For example, many studies on resource distribution tend to use positively valenced rewards, while the effort costs associated with distributing tasks are typically negatively valenced. The valence of the object to be distributed may affect participants' decisions about what distribution strategies to use. It is also likely that decisions relating to task and resource distribution are sensitive to how the object of the distribution is framed (c.f. prospect theory; Tversky & Kahneman, 1979) and may in fact lead to different default framings. It is also important for future research to compare between situations where participants must plan and decide how to distribute tasks (as in the current study) compared with situations where participants are the recipients of unfair task distributions who can only evaluate a given distribution (e.g. Blake & McAuliffe, 2011; Dawes et al., 2007).

Our findings contribute to a growing body of literature on how individuals consider rationality when engaging in joint actions. The fact that participants are willing both to incur individual motor costs to themselves (Experiment 1; see also behaviour in the incongruent condition in Török et al., 2019) and to systematically assign such costs to others (Experiment 2; see also behaviour in congruent condition in Török et al., ibid.) indicates that these sequential joint actions are represented in a way that emphasises the shared goal and action plans that maximise the joint utility. Such shared goal representations are central to theories of joint action (Gallotti & Frith, 2013; Keller, 2008; Vesper et al., 2010) and are important for facilitating interpersonal coordination. This task representation need not be prosocial, as we do not see a bias towards being fair or minimising the partner's contribution when the two path lengths are equal (in contrast with Török et al., ibid). Instead, these findings show that joint actions follow the same principle of rationality as individual actions, even if their constituent individual contributions appear to violate it.

In conclusion, the current study supports previous studies finding that individuals prioritise the efficiency of a joint task over the efficiency of their individual actions, and further demonstrates that asymmetry or inequality in the individual contributions to a joint action do not appear to be considered when distributing joint tasks. Taken in light of previous research and current theories of joint action, this indicates that people form superordinate goal representations at the level of joint actions, and that these shared goal representations are relatively insensitive to individual inefficiencies or unfair distributions of effort.

Acknowledgements

This work was supported by the European Research Council under the European Union's Seventh Framework Program (FP7/2007-2013) / ERC grant agreement no. 616072, JAXPERTISE awarded to Prof. Natalie Sebanz (Department of Cognitive Science, Central European University). The data files and experimental scripts are available from the OSF at the following link: https://osf.io/ay7n4/

Footnotes

1. 12 participants were excluded from the analysis of Experiment 1 due to low performance on the perceptual judgements task. However, 11 of these excluded subjects scored with less than 30% accuracy, indicating that they were doing the task correctly but were confused about the response mappings. The analysis of Experiment 1 reported in the main text does not include these participants, but an analysis with these 11 participants included (total $n=35$) finds the same results: participants were significantly more likely to choose the symmetrical path when it was coefficient than when the two paths were equally coefficient ($\beta=-1.58$, $SE=0.08$, $z=-19.31$, $p<.001$), and were more likely to choose the asymmetrical path when it was coefficient than when the two paths were equal ($\beta=1.65$, $SE=0.08$, $z=20.18$, $p<.001$). With this full sample there was no evidence for a bias towards either side on equally coefficient trials regardless of whether the asymmetry led to a longer path for the other ($\beta=-0.04$, $SE=0.05$, $z=-0.73$, $p=.464$), or the self ($\beta=-0.01$, $SE=0.07$, $z=-0.21$, $p=.836$).

References

Blake, P. R., & McAuliffe, K. (2011). “I had so much it didn’t seem fair”: Eight-year-olds reject two forms of inequity. *Cognition*, *120*(2), 215–224. https://doi.org/10.1016/j.cognition.2011.04.006

Brethel-Haurwitz, K. M., Stoycos, S. A., Cardinale, E. M., Huebner, B., & Marsh, A. A. (2016). Is costly punishment altruistic? Exploring rejection of unfair offers in the Ultimatum Game in real-world altruists. *Scientific Reports*, *6*(1), 1–10. https://doi.org/10.1038/srep18974

Capraro, V., & Rand, D. G. (2018). Do the Right Thing: Experimental evidence that preferences for moral behavior, rather than equity or efficiency per se, drive human prosociality. *Judgment and Decision Making*, *13*(1), 99–111. https://econpapers.repec.org/article/jdmjournl/v_3a13_3ay_3a2018_3ai_3a1_3ap_3a99-111.htm

Civai, C., Corradi-Dell’Acqua, C., Gamer, M., & Rumiati, R. I. (2010). Are irrational reactions to unfairness truly emotionally-driven? Dissociated behavioural and emotional responses in the Ultimatum Game task. *Cognition*, *114*(1), 89–95. https://doi.org/10.1016/j.cognition.2009.09.001

Clark, C. J., Baumeister, R. F., & Ditto, P. H. (2017). Making punishment palatable: Belief in free will alleviates punitive distress. *Consciousness and Cognition*, *51*, 193–211. https://doi.org/10.1016/j.concog.2017.03.010

Cohen, R. G., & Rosenbaum, D. A. (2004). Where grasps are made reveals how grasps are planned: Generation and recall of motor plans. *Experimental Brain Research*, *157*(4), 486–495. https://doi.org/10.1007/s00221-004-1862-9

Colman, A. M., Pulford, B. D., & Rose, J. (2008). Collective rationality in interactive decisions: Evidence for team reasoning. *Acta Psychologica*, *128*(2), 387–397. https://doi.org/10.1016/j.actpsy.2007.08.003

Csibra, G., Bíró, S., Koós, O., & Gergely, G. (2003). One-year-old infants use teleological representations of actions productively. *Cognitive Science*, *27*(1), 111–133. https://doi.org/10.1207/s15516709cog2701_4

Dawes, C. T., Fowler, J. H., Johnson, T., McElreath, R., & Smirnov, O. (2007). Egalitarian motives in humans. *Nature*, *446*(7137), 794–796. https://doi.org/10.1038/nature05651

Dockendorff, M., Sebanz, N., & Knoblich, G. (2019). Deviations from optimality should be an integral part of a working definition of SMC: Comment on “The body talks: Sensorimotor communication and its brain and kinematic signatures” by Pezzulo et al. *Physics of Life Reviews*, *28*, 22–23. https://doi.org/10.1016/j.plrev.2019.01.010

Fehr, E., & Gächter, S. (2002). Altruistic punishment in humans. *Nature*, *415*(6868), 137–140. https://doi.org/10.1038/415137a

Gallotti, M., & Frith, C. D. (2013). Social cognition in the we-mode. *Trends in Cognitive Sciences*, *17*(4), 160–165. https://doi.org/10.1016/j.tics.2013.02.002

Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naïve theory of rational action. *Trends in Cognitive Sciences*, *7*(7), 287–292. https://doi.org/10.1016/S1364-6613(03)00128-1

Gu, R., Yang, J., Shi, Y., Luo, Y., Luo, Y. L. L., & Cai, H. (2016). Be Strong Enough to Say No: Self-Affirmation Increases Rejection to Unfair Offers. *Frontiers in Psychology*, *7*. https://doi.org/10.3389/fpsyg.2016.01824

Keller, P. E. (2008). Joint action in music performance. In *Enacting intersubjectivity: A cognitive and social perspective on the study of interactions* (pp. 205–221). IOS Press.

Lepine, J. A., & Van Dyne, L. (2001). Peer Responses to Low Performers: An Attributional Model of Helping in the Context of Groups. *Academy of Management Review*, *26*(1), 67–84. https://doi.org/10.5465/amr.2001.4011953

Lyons, J., Hansen, S., Hurding, S., & Elliott, D. (2006). Optimizing rapid aiming behaviour: Movement kinematics depend on the cost of corrective modifications. *Experimental Brain Research*, *174*(1), 95–100. https://doi.org/10.1007/s00221-006-0426-6

Melis, A. P., Altrichter, K., & Tomasello, M. (2013). Allocation of resources to collaborators and free-riders in 3-year-olds. *Journal of Experimental Child Psychology*, *114*(2), 364–370. https://doi.org/10.1016/j.jecp.2012.08.006

Meyer, M., van der Wel, R. P., & Hunnius, S. (2013). Higher-order action planning for individual and joint object manipulations. *Experimental Brain Research*, *225*(4), 579–588.

Pezzulo, G., Donnarumma, F., & Dindo, H. (2013). Human Sensorimotor Communication: A Theory of Signaling in Online Social Interactions. *PLOS ONE*, *8*(11), e79876. https://doi.org/10.1371/journal.pone.0079876

Ray, M., de Grosbois, J., & Welsh, T. N. (2017). Index of difficulty and side of space are accommodated during the selection and planning of a joint action. *Human Movement Science*, *54*, 197–209. https://doi.org/10.1016/j.humov.2017.05.009

Ray, M., & Welsh, T. N. (2011). Response Selection During a Joint Action Task. *Journal of Motor Behavior*, *43*(4), 329–332. https://doi.org/10.1080/00222895.2011.592871

Rosenbaum, D. A., Marchak, F., Barnes, H. J., Vaughan, J., Slotta, J. D., & Jorgensen, M. J. (1990). Constraints for action selection: Overhand versus underhand grips. In M.

Jeannerod (Ed.), *Attention and performance 13: Motor representation and control* (pp. 321–342). Lawrence Erlbaum Associates, Inc.

Rosenbaum, D. A., Meulenbroek, R. J., Vaughan, J., & Jansen, C. (2001). Posture-based motion planning: Applications to grasping. *Psychological Review*, *108*(4), 709–734. https://doi.org/10.1037/0033-295x.108.4.709

Rosenbaum, D. A., Slotta, J. D., Vaughan, J., & Plamondon, R. (1991). Optimal movement selection. *Psychological Science*, *2*(2), 86–91.

Royka, A., Aboody, R., & Jara-Ettinger, J. (2018). Movement as a message: Inferring communicative intent from actions. *CogSci*, 6.

Santamaria, J. P., & Rosenbaum, D. A. (2011). Etiquette and Effort: Holding Doors for Others. *Psychological Science*, *22*(5), 584–588. https://doi.org/10.1177/0956797611406444

Schmidt, M. F. H., & Sommerville, J. A. (2011). Fairness Expectations and Altruistic Sharing in 15-Month-Old Human Infants. *PLOS ONE*, *6*(10), e23223. https://doi.org/10.1371/journal.pone.0023223

Schmitz, L., Vesper, C., Sebanz, N., & Knoblich, G. (2017). Co-representation of others' task constraints in joint action. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(8), 1480–1493. https://doi.org/10.1037/xhp0000403

Sloane, S., Baillargeon, R., & Premack, D. (2012). Do Infants Have a Sense of Fairness? *Psychological Science*, *23*(2), 196–204. https://doi.org/10.1177/0956797611422072

Southgate, V., Begus, K., Lloyd-Fox, S., di Gangi, V., & Hamilton, A. (2014). Goal representation in the infant brain. *NeuroImage*, *85*, 294–301. https://doi.org/10.1016/j.neuroimage.2013.08.043

Tabibnia, G., & Lieberman, M. D. (2007). Fairness and cooperation are rewarding: Evidence from social cognitive neuroscience. *Annals of the New York Academy of Sciences*, *1118*, 90–101. https://doi.org/10.1196/annals.1412.001

Taggar, S., & Neubert, M. J. (2008). A Cognitive (Attributions)-Emotion Model of Observer Reactions to Free-Riding Poor Performers. *Journal of Business and Psychology*, *22*(3), 167–177. https://doi.org/10.1007/s10869-008-9058-0

Török, G., Pomiechowska, B., Csibra, G., & Sebanz, N. (2019). Rationality in Joint Action: Maximizing Coefficiency in Coordination. *Psychological Science*, *30*(6), 930–941. https://doi.org/10.1177/0956797619842550

Török, G., Stanciu, O., Sebanz, N., & Csibra, G. (submitted). *Joint action planning: Co-actors minimize the aggregate individual costs of actions* [Submitted manuscript].

Tversky, A., & Kahneman, D. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291.

Vesper, C., Butterfill, S., Knoblich, G., & Sebanz, N. (2010). A minimal architecture for joint action. *Neural Networks*, *23*(8–9), 998–1003.

Vesper, C., Van Der Wel, R. P. R. D., Knoblich, G., & Sebanz, N. (2013). Are You Ready to Jump? Predictive Mechanisms in Interpersonal Coordination. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(1), 48–61. https://doi.org/10.1037/a0028066

Wang, Y., & Henderson, A. M. E. (2018). Just rewards: 17-Month-old infants expect agents to take resources according to the principles of distributive justice. *Journal of Experimental Child Psychology*, *172*, 25–40. https://doi.org/10.1016/j.jecp.2018.02.008