Joint action planning: co-actors minimize the aggregate individual costs of actions

Georgina Török (Torok_Georgina@phd.ceu.edu)
Department of Cognitive Science, Central European University, Nádor u. 9., Budapest, 1051, Hungary

Oana Stanciu (Stanciu_Oana@phd.ceu.edu)
Department of Cognitive Science, Central European University, Nádor u. 9., Budapest, 1051, Hungary

Natalie Sebanz (SebanzN@ceu.edu)
Department of Cognitive Science, Central European University, Nádor u. 9., Budapest, 1051, Hungary

Gergely Csibra (CsibraG@ceu.edu)
Department of Cognitive Science, Central European University, Nádor u. 9., Budapest, 1051, Hungary
Department of Psychological Sciences, Birkbeck, University of London, Malet Street, Bloomsbury
London WC1E 7HX, United Kingdom

Abstract

Successful cooperative activities rely on the efficient distribution of sub-tasks between co-actors. Previous research has found that people often forgo individual efficiency in favor of group-level efficiency (i.e., joint cost minimization) when planning a joint action. The present study investigated the cost computations underlying such "co-efficient" decisions: We tested the hypothesis that people compute the joint costs of a shared action sequence by summing the individual costs of their own and their co-actor's actions. We independently manipulated the parameters quantifying individual and joint action costs and tested their effects on decision-making. Participants weighed their own and their partner's costs equally to estimate the joint action costs as the sum of the two individual parameters. The results provide empirical support for computational approaches that formalize cooperation as joint utility maximization based on a sum of individual action costs.

Keywords: cooperation; joint action; efficiency; decision-making; cost computation

Introduction

Humans cooperate: in joint actions, they share goals with others, and plan and coordinate actions with their partners to achieve these goals (Butterfill, 2017; Sebanz, Bekkering, & Knoblich, 2006). Often, multiple sub-tasks contribute to an overarching joint goal, and the corresponding actions may be distributed among co-actors in many different ways with varying degrees of efficiency. Planning cooperative activities can be regarded as making a series of decisions about the actions to be performed. What principles guide people's decision-making in joint action contexts?

Do actors plan their actions to optimize efficient performance as individuals, or do they plan them with their partners in mind, to be maximally efficient as a group? Previous studies suggest that co-actors tend to maximize the joint efficiency of an action sequence by minimizing the total costs of movements when they work towards a shared goal (Kleiman-Weiner, Ho, Austerweil, Littman, & Tenenbaum, 2016; Török, Pomiechowska, Csibra, & Sebanz, 2019).

In an experimental joint action task without monetary payoffs, Török and colleagues (2019) recently observed collective utility-maximizing decisions, which they termed *co-efficient*. In their task, participants made binary decisions between two action plans to coordinate the transfer of an object with a partner. Costs in that context were operationalized as the length of movement paths to be covered on a touchscreen. One of the action options was more efficient for the decision-making actor (i.e., the initiator of the action sequence), while the other option was more efficient for her partner; the co-efficient option coincided with either. The study tested whether the participants maximized their own efficiency, and when they did not, whether the decisions followed a co-efficiency maximizing strategy or an altruistic strategy to increase the partner's individual utility (cf. Trivers, 1971). The results supported the co-efficiency hypothesis: Participants made decisions that minimized the joint, rather than the individual, action costs – even if the co-efficient solution required additional individual effort from either member of the dyad (Török et al., 2019).

The present study focuses on the computations that underlie co-efficient decisions. To minimize a dyad's costs in action planning, a decision-maker first needs to reliably estimate them (Körding & Wolpert, 2006). In joint actions, beyond their own individual action costs, people are sensitive to their partner's efforts, needs, and task difficulty (Ray, de Grosbois, & Welsh, 2017). This sensitivity is reflected in participants' increased willingness to invest effort in the joint action (Chennells & Michael, 2018), and in adjusting their actions to facilitate the partner's task (Ray et al., 2017). We hypothesize that, whenever the summed total of individual costs is calculable, it is taken as the estimate of the cost of a joint action. This proposal gains support from computational work that has formulated cooperative planning as maximizing the utility of a collective agent, in which joint utility is computationally represented as the weighted sum of the individual utilities of each agent (Kleiman-Weiner et al., 2016).

While assessing and summing individual costs may be the generic process to calculate joint costs to be minimized, depending on the context, shortcuts may also be available. For example, the task employed by Török and colleagues (2019) required actors to move an object along one of two paths to a target location. Although the movement along these paths was divided between the participants, the decision-making actor could have just planned the joint action sequence as if she had intended to complete the task alone, and then performed only the first section of the plan. Such a planning process would choose the co-efficient action option from the alternatives without requiring the planner to sum individual costs.

In the present study, we employed a novel task in which joint action costs cannot be computed without representing and summing two individual action costs. As in Török et al. (2019), action costs were estimated to be proportional to the path length of movements. However, the physical separation of paths to be taken by co-actors (see Figure 1) made it impossible to plan a single action that incorporated both paths. This feature of the task enabled us to generate, and parametrically vary, individual action costs that were statistically independent from each other. If people represent the joint costs of an action sequence as a sum of individual costs, choices between action plans should be consistent with a co-efficiency maximizing strategy that minimizes this sum. We hypothesized that, in the absence of asymmetries in cost-related uncertainty or in social hierarchy, which could potentially justify an unequal weighting of costs (Kleiman-Weiner et al., 2016), the individual costs of the actors would be weighed equally in the sum.

Our task also allowed us to investigate an alternative hypothesis related to fairness. People are often motivated to reduce payoff inequality in economic games (Dawes, Fowler, Johnson, McElreath, & Smirnov, 2007). Similarly, actors in the present study might choose action plans to minimize the difference in the action costs distributed across co-actors, even at the expense of individual or group efficiency.

Methods

Participants

Forty-eight participants (24 dyads) took part in the experiment. We analyzed the data of 40 participants (15 males, 24 females, 1 preferred not to disclose, age M = 24.21 years, SD = 4.09) after the exclusion of three dyads whose data were partially (2 dyads) or wholly (1 dyad) missing due to equipment failure. In addition, we excluded and subsequently replaced one dyad because the sampled Self and Other action costs were statistically significantly correlated in their data (see below).

Apparatus

The task was performed on a touchscreen monitor (Iiyama PROLITE 46”, resolution 1920 X 1080 pixels, horizontal sync: 31.47 – 67.5 KHz, vertical sync: 47 – 63 Hz) lying flat on a table between two participants facing each other, connected to an Apple MacBook Air computer. Stimulus presentation and data recording were controlled by a script using the Psychophysics Toolbox (Kleiner et al., 2007) in MATLAB®. Two response boxes (Black Box Toolkit Ltd.) were used to control trial onset.

Stimuli and Task

On each trial, a layout displaying the following elements was presented to the participants: (1) a thin black wall dividing the screen into two halves, corresponding to the two participants’ task areas, (2) two pairs of black target objects (two circles and two squares, 30 pixels [px] diameter) distributed between the two task areas (one of each), and (3) two black-bordered octagonal starting locations (96x96 px) with two smaller octagons inside (60x60 px, see Figure 1). The starting locations were always displayed at mid-position along the longer sides of the screen, aligned with the response box keys. At the beginning of each trial, one of the smaller octagons was orange-colored to signal which participant would initiate the joint action (playing the role of Actor 1). After 3 seconds, the color switched to green, which served as a cue for Actor 1 to start to move.

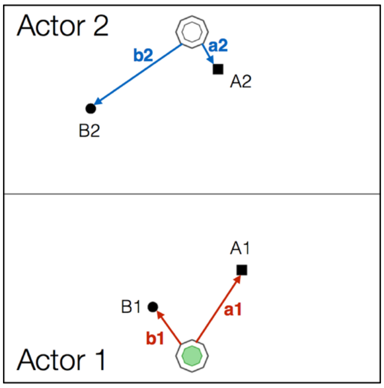

Figure 1. An example of the trials in the task. The starting locations were indicated by the octagons, and the locations for the two pairs of black target objects were generated by stochastic selection processes. The arrows and labels indicate the distances that provided the basis for cost calculations comparing the two target options, and were not visible to the participants.

In each trial, Actor 1 had to choose between the two target objects on her side and drag the chosen target back to her starting location using the green octagon. Having completed this, her partner (Actor 2) had to pick up the corresponding object from his task area and drag it back to his starting location. Thus, while both participants acted in each trial, only one of them (Actor 1) made the decision that determined the individual and joint costs incurred in a trial.

Design

Cost Disparities The cost of an action was considered to be a monotonic function of the path length that the object covered on the touchscreen when being dragged. For the sake of simplicity, here we treat the path length as the absolute cost paid for its transport. For example, in Figure 1, the cost of choosing object A1 (the square) is the distance between Actor 1's starting location and this square: a1. If Actor 1 makes her decision based on her expected cost, she should compare this cost to the cost of moving object B1: b1. The cost disparity between these actions is expressed by the difference between a1 and b1. We will call this value Self Cost Disparity, or simply Self Disparity.

If Actor 1 intends to make individually efficient decisions, she should choose A1 when the Self Cost Disparity is negative, and B1 when this value is positive. The matching individual cost disparity for Actor 2 (Other Disparity, a2 - b2) in this example is negative. Thus, picking up object B1 would be individually optimal for Actor 1, as it minimizes her Self Disparity, whereas it is the less efficient option for Actor 2.

The joint cost of an action is taken to be the summed costs of the actors. If Actor 1 chooses A1, the joint cost is a1 + a2; if she chooses B1, the joint cost is b1 + b2. Thus, the Joint Cost Disparity (or Joint Disparity) is (a1 + a2) - (b1 + b2), which is the sum of the two individual disparities (Self Disparity and Other Disparity). In the example in Figure 1, the Joint Disparity is negative, suggesting that from the dyad's perspective, picking up the square objects (A1 and A2) was associated with the shorter total path length, and as such, was the co-efficient choice. At the same time, picking up the square object pair was also individually efficient for Actor 2, but not for Actor 1. This illustrates the fact that depending on the spatial configuration of the objects in a trial, the cost-minimizing interests of each individual actor may or may not align with the interest of the dyad they are part of.

We assume that the likelihood of choosing object A1 parametrically depends on the magnitude of one or more of these disparities through a logistic link function. For example, if Actor 1 optimizes her own cost, the more negative the value of Self Disparity, the more likely it is that she will choose A1, thereby forcing Actor 2 to act on A2.

Parameter sampling Our primary aim was to investigate the independent contributions of Self Disparity and Other Disparity to the actors' decisions. To make this possible, we kept the distributions of these two factors uncorrelated across trials. To generate the locations of the target objects, we sampled Self Disparity and Other Disparity for each trial independently from the same uniform distribution (between -265 and 265 px). We then randomly selected the positions of all objects in such a way as to match these disparities. For each dyad, we generated 100 different spatial arrangements. This list was repeated twice, once per each participant acting as Actor 1, totaling 200 trials per dyad presented in pseudo-random order, with the constraint that neither of the participants be assigned the role of Actor 1 more than 3 times in a row.

The sampling process that generated object arrangements guaranteed that Self Disparity and Other Disparity were uncorrelated. As a direct consequence, Joint Disparity (the sum of the two individual disparities) had a triangular distribution and was positively correlated with both terms.

To address the alternative hypothesis regarding Fairness, we operationalized it as the difference between the asymmetries in individual paths related to object pair A and object pair B, distributed between co-actors in each trial, that is, as the difference between [abs(a1-a2)] and [abs(b1-b2)] (see Figure 1). De-correlating Self and Other Disparities from one another also de-correlates Joint Disparity from Fairness, which allowed us to estimate a model that included both of these predictors and to compare it to a single-predictor Fairness model.

Procedure

The participants were instructed to collect matching object pairs by cooperating with their partner without communicating with one another, and to complete each trial as quickly as possible. No feedback was provided about performance. Participants took on average M = 34.95 minutes (SD = 5.06 minutes) to finish the task.

To trigger the start of each trial, the participants were instructed to keep their dominant index finger on the key of their response box. First, an orange-colored octagon appeared inside one of the starting locations, which identified the participant who was required to start the trial (Actor 1). Participants were instructed to inspect the layout while the octagon was orange-colored, and to decide which target object they would pick up when prompted to move.

The octagon turned green after three seconds, signaling that Actor 1 could start to move the octagon to one of the objects. By dragging the green octagon over a black object with her index finger, the participant picked up the object and collected it by dragging it back to her starting location. Once Actor 1 collected an object, she pressed the key on her response box again to make the white octagon in front of her partner turn green. The appearance of this second green octagon cued Actor 2 to start moving to collect the matching object on their side of the screen. The trial was over when Actor 2 collected the object with the shape corresponding to the one chosen by Actor 1 (non-matching objects did not respond to dragging).

Data Analysis

To test the hypothesis that object choices would be influenced primarily by the difference between aggregate joint action costs, we used hierarchical logistic regression models in a Bayesian parameter estimation framework (Kruschke, 2015). We fitted and contrasted three models in which the probability of Actor 1 choosing object A1 was predicted in turn, by

(1) *Self Disparity*, (2) *Other Disparity*, and (3) a weighted linear combination of the *Self and Other Disparities*. In addition, we also fitted and contrasted models that attempted to predict choices by (4) Fairness or by (5) the linear combination of Fairness and Joint Disparity.

We assumed that the trial-by-trial probability of choosing A1 was Bernoulli distributed with parameter $\mu_{i|s}$, where i indexes the trial and s indexes the participant. The value of this parameter depended on a logistic function of the focal cost parameter of the model weighted by the participant's beta coefficient $\beta_{1,s}$ (for the combination models, by $\beta_{1,s}$ and $\beta_{2,s}$). We fixed the value of the intercept in the logistic equation to 0 in all models, equivalent to assuming random decisions in the absence of any action cost disparities.

The individual β coefficients were assumed to be normally distributed at the group level. We set the uninformed priors for this group-level distribution by vague hyperparameters ($\mu \sim \mathcal{N}(0, 2)$, $\sigma \sim \mathcal{U}(0.0, 0.5)$), a wide distribution around a zero effect of cost disparity. The same uninformed hyperprior was used for all cost disparities, expressing our prior expectation that participants would weigh the minimization of all costs equally. The individual and group-level posterior distributions of the beta coefficients were estimated via Markov Chain Monte Carlo simulation in JAGS, following Kruschke (2015).

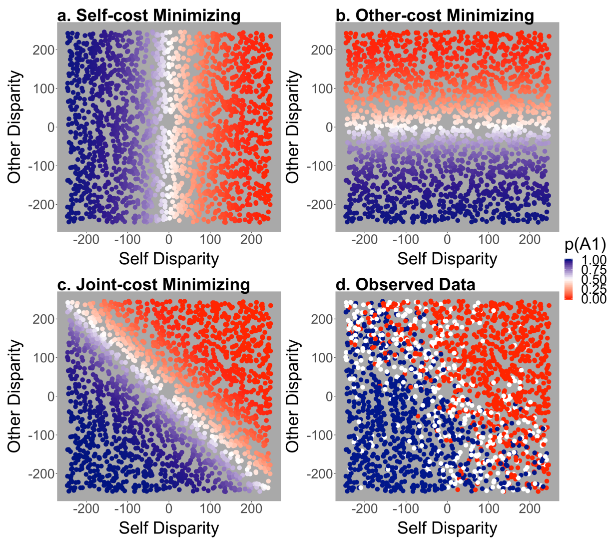

Figure 2. (a-c) Predictions for optimal responses according to Self, Other, and Joint cost-minimizing strategies (calculated assuming a Boltzmann policy with the temperature parameter fixed to $k = 50$). (d) Observed object A1 choices for all trials. Each dot in the scatter plot corresponds to the cost parameters of a trial, presented twice to a dyad (once to each participant as Actor 1). Red dots indicate that neither participant in a dyad chose object A1 in the trial; white dots: one of them chose object A1; and blue dots: both of them chose A1.

Results

Overall, the participants chose object A1 1981 times out of a total of 4000 trials. Further, at the individual level, most participants' object choices were not different from chance, suggesting that they were not biased to pick up objects based on their shape (one-sample t-test comparing the group to chance: $t(39) = 0.61$, $p = .542$, Cohen's $d = 0.10$, 95% confidence interval for proportion .50 = [.48, .51]).

Descriptive Statistics

Cost-minimization Cost disparities strongly influenced object choices, and the effects observed were a close qualitative match with the predictions of the Joint Cost-minimizing strategy (see Figure 2). Overall, participants chose the object resulting in a co-efficient action sequence 3235 times out of a total of 4000 trials (80.9%). The proportion of co-efficient choices was significantly higher than chance (.5) for each participant as well as in the whole sample ($t(39) = 33.23$, $p < .001$, Cohen's $d = 5.26$, 95% CI for proportion .81 = [.79, .83]).

Fairness We found that participants made fair choices 2071 times in the 4000 trials (51.8%). On the group level, the proportion of fair choices was statistically significantly above the chance level (one-sample t-test of individual proportions to .5: $t(39) = 2.62$, $p = .013$, Cohen's $d = 0.41$, 95% CI for proportion .52 = [.50, .53]). However, the choices were much more strongly influenced by Joint Cost-Minimizing concerns than by Fairness (see Figure 3).

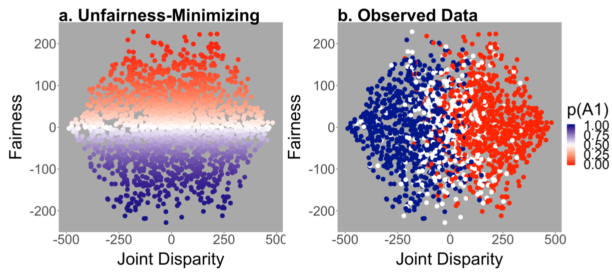

Figure 3. (a) Predictions for optimal responses according to a strategy that minimizes the unfairness of task distribution between co-actors (calculated assuming a Boltzmann policy with the temperature parameter fixed to $k = 50$). (b) Observed object A1 choices for all trials as a function of Fairness and Joint Cost Disparity.

Parameter Estimations

Models 1 and 2: Self Disparity, Other Disparity On their own, Actor 1's individual cost disparities had a significant effect on the probability of their choosing object A1. The estimated mode of the posterior for μ_β, the parameter denoting the group level weight for the cost disparity, was -.013, and its 95% highest density interval (HDI) was [-.016, -.010], which, crucially, excludes zero (Figure 4a). This means that

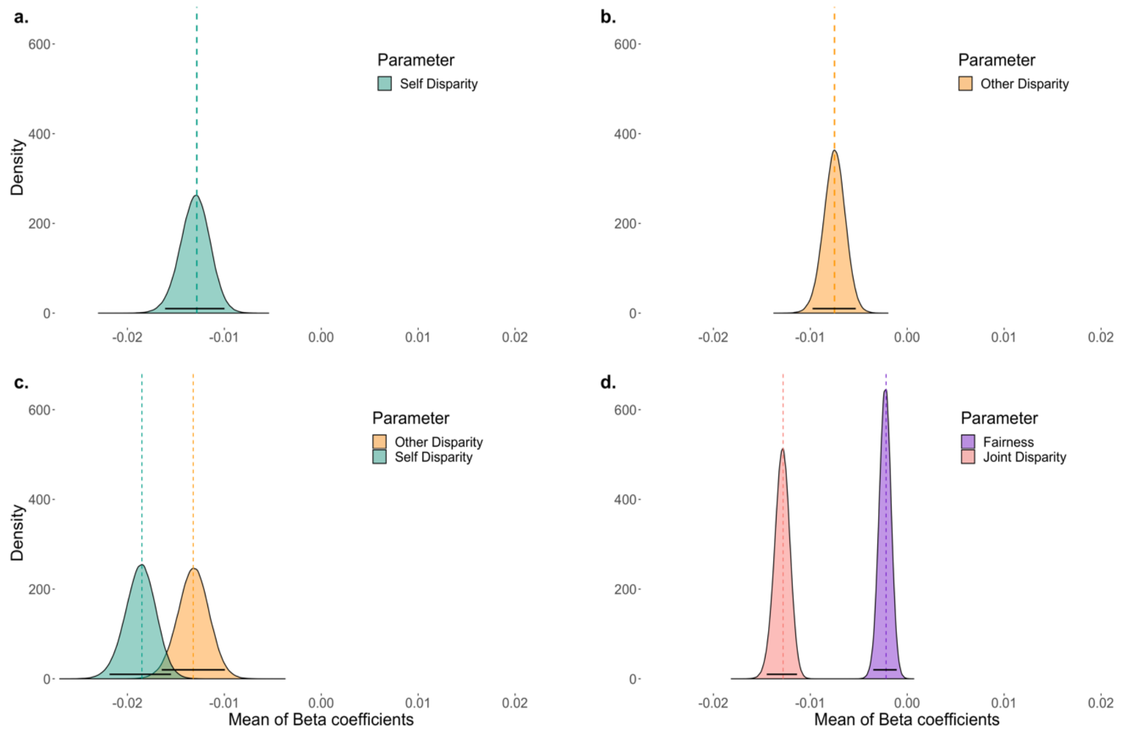

Figure 4. Posterior probability distributions of the μ_β parameters (a-b) for Self and Other Disparities in the single predictor models (Models 1 & 2) and in the (c) combination model (Model 3), and (d) for Joint Disparity + Fairness (Model 5). Dashed vertical lines indicate the Mode μ_β, the black horizontal lines represent the 95% highest density intervals of each distribution.

a one px increase in the cost disparity for Actor 1 is expected to result in the odds of an A1 choice over a B1 choice decreasing by exp(Mode μ_β) = .987. That is, for every 100 px (approx. 5.3 cm distance on the screen) increase in Self Disparity, a 73% decrease in the odds is expected.

The estimation of the μ_β parameter's posterior distribution for the Other Disparity model revealed that Actor 2's cost disparities also had a statistically significant effect on the odds of Actor 1's object A1 choices, when considered on their own. The 95% HDI did not include zero ([-.010, -.005]), with the most credible value for μ_β being -.008 (Figure 4b). Given the most credible estimates, the odds of choosing A1 over B1 were estimated to decrease by exp(Mode μ_β) = .993 for a one px increase in the Other's cost disparity. This translates to a 55% estimated decrease in the odds for every 100 px increase in Other Disparity.

The results of these two single-predictor models suggest that both individual cost disparities influenced the probability of Actor 1 choosing object A1.

Model 3: Self and Other Disparities Estimated μ_β posterior distributions for the model containing both Self and Other Disparities were similar to the two single-predictor models' (Figure 4c). The group-level means ($\mu_{\beta1}$ and $\mu_{\beta2}$) of the β_1 and β_2 coefficients for both Self and Other Disparities, respectively, were distributed below zero (Self: 95% HDI for $\mu_{\beta1}$: [-0.022, -0.016], Mode $\mu_{\beta1}$ = -0.019; Other: 95% HDI for $\mu_{\beta2}$: [-0.016, -0.010], Mode $\mu_{\beta2}$ = -0.013). Therefore, the odds of choosing A1 decreased when Self and Other Disparities increased (odds ratio for a one px increase: exp(Mode $\mu_{\beta1}$) = 0.982; exp(Mode $\mu_{\beta2}$) = 0.987). Increasing Self and Other disparities by 100 px is expected to lead to an 85% and 73% decrease in the odds of picking A1 over B1, respectively.

The 95% HDIs of the combination weights of the two cost disparities (at the group level) overlapped with one another, suggesting no difference between the magnitudes of the effects of the two parameters on decision-making. The estimated standard deviations of the two group-level posterior distributions were not different from one another, or from the standard deviations of the single-predictor models.

Models 4 & 5: "Minimizing unfairness" The 95% HDI of the posterior distribution of group-level μ_β of the single-predictor Fairness model did not include zero ([-0.002, -0.0004]), and the most credible β coefficient was Mode μ_β = -0.001. These results suggest that Fairness weakly influenced the probability of Actor 1 choosing A1. In the combination model, we found a similar effect (Figure 4d), but with a much larger effect of co-efficiency (Joint Disparity: 95% HDI for $\mu_{\beta1}$: [-0.014, -0.011], Mode $\mu_{\beta1}$ = -0.013; Fairness: 95% HDI for $\mu_{\beta2}$: [-0.004, -0.001], Mode $\mu_{\beta2}$ = -0.002). The odds of choosing A1 over B1 decreased by exp(Mode $\mu_{\beta1}$) = 0.987 with every one px increase of the Joint Disparity, whereas the odds decreased by exp(Mode $\mu_{\beta2}$) = 0.998 with a one px increase in unfairness of the cost distribution. Increasing Joint Disparity and the Fairness asymmetries by 100 px is expected to lead to a 73% and 18% decrease in the odds of picking A1 over B1, respectively. The estimated standard deviations for the two parameters' posterior distributions were not different from one another.

Model Comparison

To compare model fits to the data, we extracted Deviance Information Criterion (DIC) for each model. We found that the Self + Other Disparities (i.e., Joint Disparity) model (Model 3) fit the choices the best (DIC = 2467), while the separate Self (DIC = 3810) and Other Disparities (DIC = 4575) models (Models 1 & 2) fit them less well. The Fairness (Model 4) was a particularly ill-fitting model (DIC = 5539). A combination model of Joint Disparity + Fairness (Model 5) predicted the data second-best (DIC = 3205).

Exploratory Analyses

We conducted explorative analyses to see if participants learned to be co-efficient over time. The task was divided into eight blocks of 25 trials, between which participants were offered the chance to take short breaks. To investigate a learning effect, we calculated the block-wise ratios of co-efficient object choices for each participant. We found no effect of time: the ratios of co-efficient choices were significantly higher than chance from Block 1, and we observed no statistically significant changes over time (Kruskal-Wallis test: $\chi^2 = 8.11$, $df = 7$, $p = .323$).

The overall ratio of co-efficient choices we found suggests that in about 1 out of 5 trials, the decisions did not minimize the joint action costs. To examine the factors driving decision-making in these cases, we analyzed the subset of trials with jointly sub-efficient decisions (765 trials, Mdn = 19.5 trials per participant, SD = 5.88). We found that in total, 23.5% of the participants' jointly sub-efficient choices minimized Self costs only, while 15.3% of them minimized Other costs only, and 9% of the choices only ensured Fairness. The remaining trials were, by design, characterized by overlaps between predicted optimal solutions for Self cost minimization, Other cost minimization and Fairness maximization: Choices in 30.6% of the jointly sub-efficient trials minimized both Self costs and unfairness, whereas in 16.2% of the trials, they minimized Other costs and unfairness. Due to the overlaps in predictions, we cannot conclude which factors determined these choices. Finally, in 5.4% of the jointly sub-efficient trials, participants' choices were inefficient and unfair to everyone (1% of the full sample). To sum up, the pattern of non-co-efficient choices suggests that minimizing Self costs was a strong secondary decision strategy, followed by some considerations for ensuring Fairness.

Additional Experiments

In two further experiments, we tested if either self- or other-cost minimization had an additional effect on choices beyond joint cost minimization. Joint Disparity was de-correlated from Self Disparity (Experiment 2) and from Other Disparity (Experiment 3). Both additional experiments adopted the same task as Experiment 1 and were indistinguishable from participants' perspective.

In the Self Disparity-only model in Experiment 2, the beta parameter did not differ from zero, indicating no effect of self-cost minimization. However, Joint Disparity as the predictor resulted in a non-zero modal beta value, comparable in magnitude to that of Experiment 1. Adding both predictors to the model did not change the results: Joint Disparity alone or with Self Disparity had a similar fit to the data.

In Experiment 3, the Other Disparity model's estimates were slightly *above* zero, indicating Actor 1's tendency to decide against minimizing Actor 2's costs. However, Joint Disparity as the sole predictor resulted in estimates that were similar to Experiments 1 and 2. Combining the two predictors revealed the same pattern of results as the single-predictor models: negative weighting for Joint Disparity, a positive one for Other Disparity. The combined model fit the data best.

Discussion

The current study addressed the computations that underlie joint cost-minimizing decisions in joint action planning. We tested the hypothesis that co-actors represent the dyad's joint costs as a sum of the members' individual costs.

In a joint object matching task, participants made binary decisions between action plans with different associated costs. Costs were operationalized as the length of paths to be taken to collect objects on a touchscreen, and the disparities between the costs of the two action options available to each actor were hypothesized to play causal roles in the decision-making process.

We independently manipulated Self and Other Cost Disparities on a trial-by-trial basis, and tested three logistic regression models that included each cost parameter and their combination. The results of the parameter estimations suggested that both individual cost disparities influenced the participants' decisions. We found the predicted negative relationship between an increase in the costs related to picking up an object and the probability of choosing it. The overlap for Self and Other Disparities and the parameter estimates of the combination model suggested that, while Self Disparity has a numerically larger effect, the size of this effect hardly differed between the two individual costs. This supports the hypothesis that on the group level, an additive combination of the two individual costs was minimized by the first actor's decisions. This result was confirmed by two additional experiments.

We also tested an alternative hypothesis according to which action decisions are determined by the minimization of unfairness in the distribution of individual action costs. We found that although fairness did have an effect on decisions (~52% of choices were "fair"), this effect was much smaller than that of the joint costs (~81% of choices were co-efficient). Joint cost minimization was a better predictor of decision-making than minimizing unfairness in task distribution.

It is worth noting, too, that in the minority of trials where the participants' decisions were guided by strategies other than joint cost minimization (~19%), they primarily tended to minimize their own individual action costs.

Regardless of the particular statistical dependencies between the cost parameters in the three experiments, models that predicted object choices based on the equally weighted sum of the co-actors' individual action costs predicted decisions with the highest accuracy. This is consistent with the way Kleiman-Weiner et al. (2016) operationalized cooperative planning, and confirms that, as long as the individual costs can be estimated on the same scale (i.e., as proportional to distance in our case), joint costs are calculated as the sum of the individual costs in joint action planning.

Future research should address the generalizability of this cost computation mechanism to different types of effort and different forms of joint action. Here, costs were operationalized such that they strongly correlated with the time required to complete a movement (path lengths). Time may serve as a common currency for comparing costs like physical and mental effort (Potts, Pastel, & Rosenbaum, 2018), but we propose that the additive cost estimation mechanism should also be applicable in cases where time is not strongly correlated with costs. This proposal could be explored by operationalizing effort as fixed-duration movement in a viscous force field.

Since our task consisted of many repeated interactions, we cannot make strong claims about the scope of the cost estimation. The participants could have estimated costs locally in each trial, or over the whole experimental session, as trial-by-trial joint cost minimization also minimized the total time spent in the lab. These hypotheses could be tested by fixing the total duration of an experimental session. Similarly, manipulating the payoff structure of the task could reveal the boundary conditions of joint cost minimization. Asymmetries in co-actors' benefits from the task would influence the extent to which they prioritize joint over individual action costs.

Investigating the factors that might modulate how individual costs are weighed in decision-making (e.g., social hierarchies, motor skill differences) will help us achieve a fuller understanding of the computations that people employ in cooperative action planning. As a first step toward this goal, the present study provides clear evidence for an additive cost computation that enables efficient coordination in joint action.

Acknowledgments

This work was supported by the European Research Council under the European Union's Seventh Framework Program (FP7/2007-2013) / ERC grant agreements n° 616072 (JAXPERTISE) and n° 742231 (PARTNERS). We thank Dávid Csűrös, Vanda Derzsi, and Fruzsina Kollányi for their assistance in data collection.

References

Butterfill, S. (2017). Joint action: A minimal account. In J. Kiverstein (Ed.), *Routledge Handbook of Philosophy of the Social Mind* (pp. 357-369). New York, NY: Routledge.

Chennells, M., & Michael, J. (2018). Effort and performance in a cooperative activity are boosted by perception of a partner's effort. *Scientific Reports*, *8*.

Dawes, C. T., Fowler, J. H., Johnson, T., McElreath, R., & Smirnov, O. (2007). Egalitarian motives in humans. *Nature*, *446*(7137), 794.

Kleiman-Weiner, M., Ho, M. K., Austerweil, J. L., Littman, M. L., & Tenenbaum, J. B. (2016). Coordinate to cooperate or compete: abstract goals and joint intentions in social interaction. *Proceedings of the 38th Annual Conference of the Cognitive Science Society*.

Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception*, *36*(14) 1-16.

Körding, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, *10*(7), 319-326.

Kruschke, J. (2015). *Doing Bayesian Data Analysis: A Tutorial with R, JAGS, and Stan* (2nd ed.). Academic Press.

Potts, C. A., Pastel, S., & Rosenbaum, D. A. (2018). How are cognitive and physical difficulty compared?. *Attention, Perception, & Psychophysics*, *80*(2), 500-511.

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.

Sebanz, N., Bekkering, H., & Knoblich, G. (2006). Joint action: bodies and minds moving together. *Trends in Cognitive Sciences*, *10*(2), 70-76.

Török, G., Pomiechowska, B., Csibra, G., & Sebanz, N. (2019). Rationality in joint action: maximizing coefficiency in coordination. *Psychological Science*, *30*(6), 930-941.

Trivers, R. L. (1971). The evolution of reciprocal altruism. *The Quarterly Review of Biology*, *46*(1), 35-57.