

Crazy for you! Understanding Utility in Joint Actions

Arianna Curioni (CurioniA@ceu.edu)

Department of Cognitive Science, Central European University.
1051 Budapest (Hungary)

Pavel Voinov

Department of Cognitive Science, Central European University.
1051 Budapest (Hungary)

Matthias Allritz

School of Psychology & Neuroscience, University of St Andrews, St. Andrews, Fife, KY16 9JU, UK

Josep Call

School of Psychology & Neuroscience, University of St Andrews, St. Andrews, Fife, KY16 9JU, UK

Gunther Klaus Knoblich

Department of Cognitive Science, Central European University.
1051 Budapest (Hungary)

Abstract

Predicting others' actions and inferring preferences from their choices is indispensable for successfully navigating social environments. Yet, the cognitive tools agents employ for prediction and decision may differ when involved in social interactions. When pursuing a goal individually, humans maximize utility by minimizing costs, while when engaged in joint actions utility maximization might not be the only heuristic in place. We investigate if human adults represent costs and rewards of joint vs. individual actions, and how do they decide whether to engage in a joint action. We test participants' decisions when solving a task alone or together with a partner as a function of the cost of coordination. Our results show that human adults decide based on a preference for joint actions, despite engaging in coordination reduces their individual utility. We discuss a framework for decision-making which accounts for cognitive heuristics and preferences for joint actions characterizing agents' cooperative behavior.

Keywords: utility; joint action; decision; coordination.

Introduction

A key challenge in social cognition research is to understand how the cognitive processes and heuristics driving individual goal directed behavior are modulated when an individual coordinates her actions to interact with another.

Dominant theories of action propose that principles of optimization and utility guide both the comprehension and the execution of actions and motor plans in humans (Wolpert, 2011). Agents, when deciding to act, use information about goals and relative costs of action implementation to select (close to) optimal action plans (Todorov, 2004). Interestingly, such principles not only guide actions' planning, but also sustain our ability to make sense of actions we observe. In the domain of developmental science, it has been argued that infants make sense of others' actions by applying assumptions of rationality and utility (Gergely &

Csibra, 2003; Jara-Ettinger et al., 2016). For example, infants are able to predict an agent's goal by assuming she would achieve it by minimizing her action costs (Gergely et al., 1995). Moreover, infants are able to infer agents' preferences based on the costs that they are willing to incur to pursue different goals (Liu et al., 2017). Conversely, when agents' choices and actions violate the assumptions of rationality, infants do not interpret them as goal directed (Hernik & Southgate, 2012; Southgate & Csibra, 2009).

The question we address is whether human adults use such utility and optimality principles when acting in social contexts. A particularly interesting case of individual actions performed in social contexts are joint actions, i.e. actions performed by two or more agents that are coordinated in space and time to bring about a change in the environment (Sebanz & Knoblich, 2006). Importantly, to successfully engage in joint actions, agents need to represent, plan, and prepare instrumental goal directed actions within the context of interpersonal coordination. Empirical evidence from a vast literature on interpersonal coordination shows that coordinating with others is more costly than performing the same actions individually. This is true whether the cost is assessed in the cognitive or motor domain. In fact, when coordinating with others, agents incur costs of representing complementary action plans (Neuman-Norlund et al., 2007; Sacheli et al., 2015) and multiple perspectives (Freundlieb et al., 2016), predicting others' actions (Kourtis et al., 2013), and monitoring more than one contribution to the action outcome (Loher et al., 2013). Moreover, when instrumental actions are performed in coordination tasks, agents would incur in extra motor costs and exploit violations of optimality as information tools for communicating intentions (Cavallo et al., 2016), disambiguating goals (Candidi et al., 2015), and facilitating synchrony (Vesper et al., 2014).

Do humans represent and weigh the relative costs and rewards of coordinated actions vs. the costs and rewards of

individual actions? How do they decide whether to engage in a joint action or individual action to achieve a particular goal?

A prominent cognitive model (i.e. Naive Utility Calculus model, Jara-Ettinger, 2016) formulates predictions about agents' goals and choices based on three fundamental assumptions: (a) agents are rational planners, (b) agents have priors over the costs of different actions among which they can choose, (c) incurring costs is indicative of a preference (higher reward) for a given goal. Here we take advantage of the structure of such model – that predicts the choice of an action based on its observable cost and infers a preference for such action based on the cost incurred - to investigate what cost-benefit computation adults apply to joint actions. This allows us to investigate if agents prefer joint action over individual action based on additional instrumental costs they are willing to incur to perform a task jointly. Any violation of individual utility (assumption (a)), defined in terms of costs and benefits of actions performed when agents are alone, may be informative of how agents represent and weigh the costs of actions when they are performed in coordination (b), and allow to infer the benefit, therefore preference for the goal of coordination (c).

In the current study we investigate adults' decisions when faced with two alternative means to solving a task: either by performing individual actions or a joint action with a partner. By manipulating the relative costs and rewards associated with individual and joint actions, we are able to exactly quantify what, if any, additional costs adults are willing to incur in acting together instead of performing the task on their own. Results from three experiments indicate that human adults show a preference for joint actions over individual actions, even when choosing joint actions reduces their individual utility.

Experiment 1

Using a touchscreen-based game, we measured adults' decisions whether to perform a joint action or an individual action, as a function of how costly it is to solve the task. Throughout the experiment, at each trial, participants' goal was to collect as many points as possible by clearing (by tapping) the touch-screen area of the 2D items (boxes) displayed on it (See Figure 1a). In 50% of trials, one of the participants, henceforth the decision-maker, had to decide whether to solve the task *alone* or *together*. For the decision-maker, solving the task *alone* implied clearing all boxes presented on the screen and gaining one point for each cleared box (e.g. clearing 8 boxes means collecting 8 points). Solving the task *together* implied splitting the boxes with the partner, therefore halving the number of actions (and number of points gained, e.g. in an 8 boxes trials, decision-maker clears 4 boxes and collects 4 points, as does the partner (See Figure 1c)). Crucially, to clear the boxes, the decision-maker and the partner had to tap on the same box at the same time (within a pre-specified time-window). If the taps were not synchronous enough, or did not land on the same box, the box would not clear from the screen and have to be tapped again.

Within a utility maximization framework, agents should always choose the task mode that maximizes the reward (number of points) and minimizes the costs (time to completion). Because solving the task *together* implies halving the costs of the individual (i.e. number of actions), but also halving the points collected at the given trial (reward), choosing to perform the task together would only be rational if participants were (at least) twice as fast at completing a trial with a joint action as completing it alone. Any additional time incurred in *together* trials would thus represent a violation of the utility maximization principle and suggest a preference for joint action over individual action (See Figure 1b).

Methods

Participants. We recruited English-speaking participants who reported no history of neurological impairments or diagnoses, and normal or corrected-to-normal vision. We recruited 40 participants in total (15 F; Average age= 26.06 y +/- 4.47).

Apparatus and Stimuli. The task was performed on an Iiyama 46" PROLITE TF4637MSC-B2AG touch-screen set (1600 × 900 pixels resolution). The screen area was vertically divided into two equal halves, corresponding to the two participants' positions during the experiment. Colored squares ("boxes", 7.8 × 7.8 cm) were displayed in a grid arrangement, equidistant from each other. Boxes were colored blue during the trial (when they were active) and green during preview (see procedure). The experimental script, trial randomization and participants' responses were controlled using MatLab 16b software.

Procedure. Participants came into the testing room in pairs and stood in front of a touchscreen that was lying flat on a table (see Figure 1). The roles of decision-maker and partner were randomly assigned. As no hypotheses about the effect of gender on the current research questions were formulated, pairs of participants were randomly composed without controlling for gender. The task for the decision-maker was to clear as many boxes as she could within 20 mins. Unknown to the participants, the number of trials was in fact fixed and the time limit was set for longer than the actual time required to finish the task. To highlight the relevance of time during the task, we installed an electronic countdown clock on the desk in front of the participants. The screen side (left-right) occupied by the decision-maker was counter-balanced across participants. No specific instructions were given as to what finger participants had to use to clear boxes on the screen, but from piloting we observed that everyone converged on using the index finger of their dominant hand. We ensured that participants could comfortably move their hands without spatial constraints, regardless of what hand participants used to complete the task. Three types of trials were presented in a random order: 45 No-Choice Alone trials, 45 No-Choice Together trials, and 90 Choice trials. At each trial 4, 8, or 12 boxes were displayed on the screen and had to be cleared. Both Choice and No-Choice trials included 30 trials of each numerosity level (with 15 Alone and 15 Together trials per




Figure 1. A. Experimental apparatus and set up. Participants’ goal at each trial is to collect as many points as possible by clearing the touch-screen area of the 2D boxes displayed on it. Decision maker (DM) and helper are standing next to each other in front of the screen. The DM has to decide whether to clear them alone or jointly with the helper. **B.** In *Alone* trials, the DM has to tap on each box and collects the same amount of points as the boxes she clears. In *Together* trials, DM and partner have to tap on the same box at the same time. The DM halves the number of actions he performs but also the points gained, e.g. in an 8 boxes trials, DM clears 4 boxes and collects 4 points. **C.** DM’s utility is measured in time to complete the trial (time from the first touch to the touch that cleared the last box on the screen): to maximize it, the DM should maximize the points per trial and minimize the time per action. Solving the task *Together* is rational only if participants are twice as fast at completing the trial compared to *Alone*. Any additional time in *Together* trials is a reduction of DM’s utility.

level). The rationale for having a fixed number of No-Choice Together and No-Choice Alone trials was to guarantee that each participant had a comparable experience with both task modes, independent of her decisions in Choice trials. Before proceeding to the experiment, participants were introduced to four practice trials, where they were familiarized with the procedure, different trial types, and variations in stimuli quantities.

At the beginning of each trial the decision-maker was presented with a preview of the upcoming trial configuration. The preview was followed by a display of two buttons (‘alone’ and ‘together’). By pressing one of the two, the decision-maker started the trial in the modality she selected and boxes changed color to signal they were now active. One of the buttons was inactive in No-Choice trials (but visually identical), so that no choice was available to the participant. Above the menu a text box reported the total number of cleared boxes (decision-maker’s score), which was updated after each trial.

In *Alone* trials boxes appeared only on the decision-maker’s side of the screen and could be cleared with a single tap of each item. In *Together* trials half of the boxes appeared on the decision-maker’s side and half of them appeared on the partner’s side of the screen. In this condition, boxes on the two halves of the screen had to be cleared simultaneously by means of a synchronous touch on the two corresponding squares (that is, those in the same row and column position)

by the two participants. By choosing *Together*, decision-makers scored points only for the boxes cleared on their side of the screen. A tolerance synchrony window of 300 ms was chosen based on the minimum interval between two touch events the touchscreen could reliably register. No time constraint was put on the time to complete a trial, although we instructed participants to be as fast as possible. Participants were also instructed not to communicate with each other, neither verbally or otherwise, and to look at the touchscreen in front of them, to avoid eye contact.

Data Analyses. The primary dependent variable was the proportion of *Together* choices, that is, the proportion of trials where individuals chose to perform a joint action. The proportion of *Together* choices was tested against 0.5 (chance level) by means of a Binomial Test.

The second dependent variable was the average Trial Time for No-Choice trials, defined as the time to complete the trial calculated from the first touch to the touch that cleared the last box. Trials that were above or below 3 standard deviations from the sample mean, calculated across conditions for each numerosity level, were excluded from the analysis (1.5% of all trials). We performed a 2 x 3 repeated measures analysis of variance (rANOVA) with Task (*Alone*, *Together*) and Number of Actions (4, 8, 12) as within subjects factors. For all analyses, the significance level was set to an α level of 0.05. Significant interactions and main effects were

analysed by Tukey post hoc tests. Data were analysed in JASP v.0.10.

Results

Results from the Binomial test show that the proportion of Together choices over Alone choices (counts = 1371/1800 observations, 0.76) was significantly larger than chance level ($p < 0.001$, CI = 0.74 lower, 0.78 upper). This indicates that participants decided to solve the task performing joint actions more frequently than individual actions (Figure 2).

As the Mauchly's test indicated a violation of the assumption of sphericity, we report Greenhouse-Geisser corrected results of the rANOVA. The results of the rANOVA on Trial Time show a significant main effect of Task ($F(1,19) = 7.54$, $p = 0.013$, $\eta^2_p = 0.28$), where participants were faster at solving the task Alone (average trial time: 1.6 sec, $sd = 0.25$) compared to Together (average trial time: 1.92 sec, 0.69), a significant main effect of Number of Actions ($F(2,38) = 288.86$, $p < 0.001$, $\eta^2_p = 0.93$), where trial duration was larger the more actions participants performed to complete the trial. The significant Task x Number of Actions interaction ($F(2,38) = 3.45$, $p = 0.04$, $\eta^2_p = 0.154$) shows that the more actions to perform per trial, the larger the difference between solving the task Alone or Together. As shown by significant post hoc tests, participants were significantly slower Together than Alone in trials with 12 boxes to clear ($p = 0.018$).

Experiment 2

In order to claim that decision-makers in Experiment 1 decided to solve the task jointly despite the higher costs, we had to demonstrate, first, that individuals were sensitive to maximizing their utility in the absence of an interactive partner (i.e. they were sensitive to the cost-benefit manipulation we introduced); and second, that the preference for joint action, revealed in Experiment 1, originated in a preference for social coordination, and not for coordinated actions per se. Experiment 2 provided a non-social control for Experiment 1. In this experiment participants always performed the task alone but chose between completing trials either uni-manually or bi-manually. It is well established that the coordinative processes regulating *intrapersonal* and *interpersonal* action coordination share striking similarities across action domains. This is particularly true for hand and limb coordination (Schmidt et al., 2008; Ramenzoni et al., 2011). The unimanual/bimanual task modes are therefore conceptually and motorically equivalent to the Alone/Together modes of Experiment 1.

Methods

Participants. We recruited English-speaking participants who reported no history of neurological impairments or diagnoses, and normal or corrected-to-normal vision. We

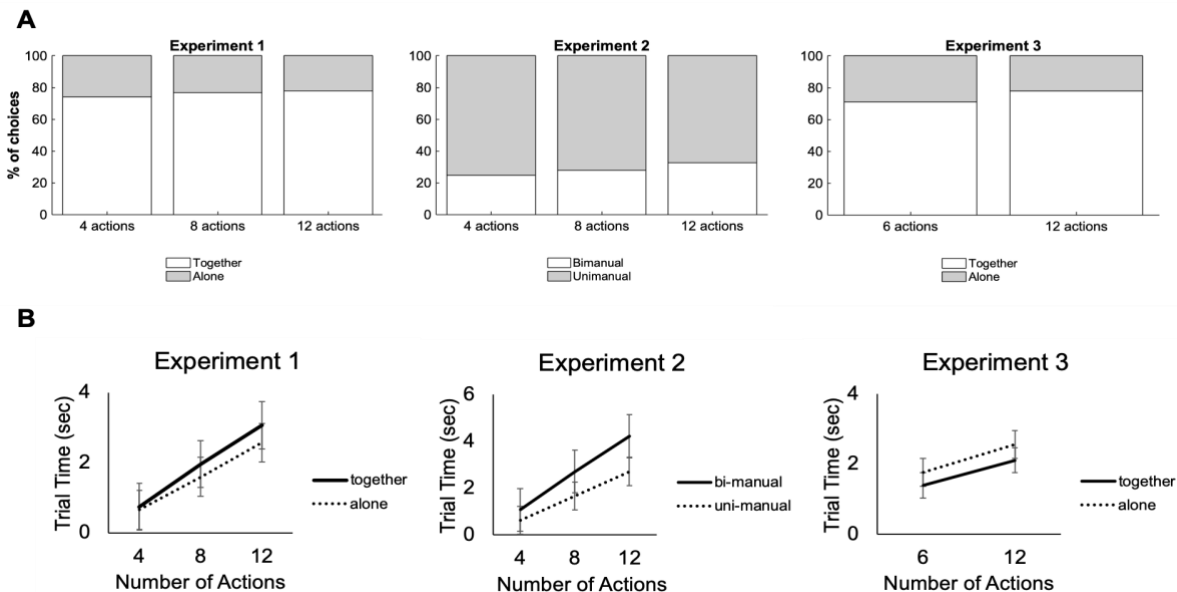


Figure 2. A. Samples' average frequencies of Together/Alone choices in Experiment 1, 3 and Bi/Unimanual choices in Experiment 2, plotted separately for each level of Number of Actions factor. B. Samples' average Trial Time in Experiment 1, 2 and 3 plotted separately for each level of Number of Actions factor. Error bars represent S.E.M.

recruited 20 participants in total (8 F; Average age= 25.8 y +/- 5.17).

Apparatus and Stimuli. Same apparatus and stimuli of Experiment 1.

Procedure. Procedure was identical to Experiment 1, with the following differences. The single participant (decision-maker) chose between two buttons ('1 hand' and '2 hands'). In Uni-manual (1 hand) trials, boxes appeared only on one side of the screen (where the decision-maker was instructed to stand, counterbalanced across participants) and could be cleared with a single tap of each item. In Bi-manual (2 hands) trials, a half of the boxes appeared on the decision-maker's side while the other half appeared on the other side of the screen. In this condition, boxes on the two halves of the screen had to be cleared simultaneously by means of a synchronous touch on the two corresponding squares by the participant with her two hands. By choosing "2 hands", decision-makers scored only the boxes they cleared on their side of the screen. This choice was therefore rational only if participants were able to complete a trial bi-manually (at least) two times faster compared to uni-manually, and without mistakes.

Data Analyses. Data analyses and trial exclusions were identical to Experiment 1 (exceeding +/- 3 SD ;1.5% of the total trial number).

Results

Results from the Binomial test show that the proportion of Bi-manual choices over Uni-manual choices (counts = 511/1800 observations, 0.28) was significantly smaller than chance level ($p < 0.001$, CI = 0.26 lower, 0.30 upper). This indicates that participants decided to solve the task uni-manually more frequently than bi-manually (Figure 2).

As the Mauchly's test indicated a violation of the assumption of sphericity, we report Greenhouse-Geisser corrected results of the rANOVA. The results of the rANOVA on Trial Time show a significant main effect of Task ($F(1,19) = 22.10$, $p < 0.001$), $\eta^2_p = 0.538$), where participants were faster at solving the task Uni-manually (average trial time: 1.66 sec, $sd = 0.25$) compared to Bi-manually (average trial time: 2.67 sec, $sd = 1.0$), a significant main effect of Number of Actions ($F(2,38) = 267.73$, $p < 0.001$, $\eta^2_p = 0.93$), where trial duration was larger the more actions participants performed to complete the trial. The significant Task x Number of Actions interaction ($F(2,38) = 14.35$, $p < 0.001$, $\eta^2_p = 0.43$) shows that the more actions to perform per trial, the larger the difference between solving the task Uni-manually or Bi-manually. As shown by significant post hoc tests, participants were significantly slower Bi-manually than Uni-manually in 8 boxes ($p = 0.002$) and 12 boxes trials ($p < 0.001$).

Experiment 3

The main goal of Experiment 3 was to test the robustness of the preference for joint action we observed in Experiment 1. We therefore modified task instructions to stress the

importance for decision-makers of maximizing the total score. This modification was aimed at highlighting for participants that the best strategy to finish the experiment faster was to choose the task mode that maximized the score (number of boxes) per trial. We also introduced two levels of task difficulty by manipulating the size of the boxes to clear on the screen. This manipulation was aimed at making the cost of clearing boxes from the screen more salient for participants, as aiming and tapping/coordinates is harder on a smaller surface (Fitts Law, 1954).

Methods

Participants. We recruited English-speaking participants who reported no history of neurological impairments or diagnoses, and normal or corrected-to-normal vision. We recruited 40 participants in total (36 F; Average age= 23.13 y +/- 2.95).

Apparatus and Stimuli. The experiment was conducted on a 43" Iiyama PROLITE TF4338MSC-B1AG touchscreen. Stimuli were identical to Experiments 1 and 2, except the boxes appeared in two sizes: big (7.3 x 7.3 cm) or small (2.2 x 2.2 cm). Synchrony threshold in Together trials was lowered to 250 ms, as the touchscreen allowed to reliably register two consecutive touches above that window.

Procedure. Procedure was identical to Experiment 1 except the following differences. Participants were given a target number of boxes to clear (1500) without a time limit. In reality, there was a fixed number of trials and the target number was beyond the possible amount the participants could collect. In addition to the trial type factor (No-Choice Alone, No-Choice Together, and Choice) and Number of Actions factor (which had only two levels: 6 and 12), there was a third factor of Box Size with two levels: Big and Small boxes. There were 20 trials of each Number x Size combination in the Choice condition (80 total), and 10 trials of each combination in Alone and Together No-Choice trials (80 for all No-Choice trials combined).

Before proceeding to the experiment, participants were introduced to four practice trials, where they were familiarized with the procedure, different trial types, and variations in stimuli size and quantities.

Data Analyses. Data analyses and trial exclusions were identical to Experiment 1 (exceeding +/- 3 SD; 2.6% of the total trial number). The repeated measures analysis of variance (rANOVA) on Trial Time had Task (Alone-Together), Number of Actions (6,12) and Boxes size (small, big) as within subjects' factors.

Results

Results from the Binomial test show that the proportion of Together choices over Alone choices (counts = 1192/1600 observations, 0.74) was significantly larger than chance level ($p < 0.001$, CI = 0.72 lower, 0.76 upper). This indicates that participants decided to solve the task performing joint actions more frequently than individual actions (Figure 2).

The results of the rANOVA on Trial Time show a significant main effect of Task ($F(1,19) = 23.7$, $p < 0.001$, η^2

$p = 0.55$), where participants were faster at solving the task Together (average trial time: 1.75 sec, $sd = 0.45$) compared to Alone (average trial time: 2.1, $sd = 0.27$), a significant main effect of Number of Actions ($F(1,19) = 593.41$, $p < 0.001$, $\eta^2_p = 0.96$), where trial duration was larger the more actions participants performed to complete the trial, and a significant main effect of Box Size ($F(1,19) = 163.19$, $p < 0.001$, $\eta^2_p = 0.89$), where participants were faster at completing the trial when boxes were bigger (as predicted by Fitts law). The significant Task x Number of Actions interaction ($F(1,19) = 6.96$, $p = 0.016$, $\eta^2_p = 0.268$) shows that the more actions to perform per trial, the larger the difference between solving the task Alone and Together. As shown by the significant post hoc test, participants were significantly slower Alone than Together in 12 boxes trials ($p < 0.001$). Finally, the significant Box Size x Number of Actions interaction shows that participants were slower at completing the trial when boxes were smaller, and the trial required more actions to be completed ($p < 0.001$).

Discussion

We investigated how people decide whether to solve a task together with a partner or alone. In many real-life scenarios, joint actions are the most efficient way to solve tasks, as coordinating with a partner allows individuals to achieve solutions that they would not be able to achieve individually. In such situations, the choice of coordinating with a partner satisfies both the goal of utility maximization for individuals and the drive to be pro-social and engage with conspecifics (Boyd & Richerson, 2009; Tomasello et al., 2012; Michael et al., 2016). However, not every task can be approached with such clear-cut solutions in terms of individual utility maximization. At times, as we show, individuals decide to perform joint actions even when performing them alone represents a more efficient alternative. Such choices pose an interesting question for social cognition theories that model agents' choices and preferences based on intuitive (or naïve) utility calculus (Jara-Ettinger, 2016). In fact, in order to account for these choices, we would need to assume that either individuals disregard (or downplay) individual utility in favour of engaging in joint actions, or the utility computation of joint actions needs to integrate more than (just) the observable costs and rewards associated with individual goal achievement.

The current results offer an initial answer to this question: human adults, when faced with the choice of solving a task together or alone, show a strong preference for joint actions (Experiment 1 and 3). This is the case despite the significant costs that coordination adds to the task, as demonstrated by the large performance differences between individual and joint trials. In the present task the computational costs of coordinating with others (Neuman-Norlund, 2007; Kourtis, 2013; Loher, 2013) resulted in observable behavioural costs, i.e. the time to complete the coordination task. The same costs, when experienced in a solo task that agents could solve either uni-manually or bi-manually (Experiment 2), were sufficient to enable individuals to choose the more efficient

way of performing the task (uni-manual). Moreover, when agents were explicitly instructed to maximize their score (Experiment 3), they still chose to complete trials in the modality that did not maximize their score (together) but did so more efficiently, i.e. by improving their coordination performance. Importantly, participants failed to achieve a performance level that rationally justifies the choice to engage in joint action (they were still not twice as fast in Together compared to Alone trials).

Individuals were clearly capable and driven to choose the most efficient means to solve the task when acting alone, as reflected by their decisions to avoid the costlier task mode (bi-manual coordination) in favour of the uni-manual solution – that ensured the highest score in the shortest time. However, when given the opportunity, participants preferred to work together in coordination with a partner and were willing to incur costs that reduce their (observable) individual utility. Within a utility maximization framework, this implies that individuals have assigned a certain preference to joint actions that justified the cost incurred.

This preference for joint actions may have resulted from the mere presence of a potential interaction partner in the room. Participants were instructed not to communicate or look at each other during the experimental session but were nevertheless acting on the experimental apparatus next to each other. A related possibility is that participants felt a need to manage their reputation in order to comply with the socially desirable behaviour of engaging with others when possible. Reputation building in social interactions is known to reduce the future risk of social exclusion (Nowak et al., 2005). Therefore, individuals may have been influenced by a risk aversion strategy when deciding to engage with the partner at this task. Being it a cost aversive or a reward seeking strategy, in either case reputation building is a likely candidate that could affect the utility calculus underlying the choice between joint and individual actions.

Further experiments, where individuals are interacting not in physical proximity, can address more directly the influence of such variables on our preference for engaging in coordinated activities despite their instrumental extra cost.

These results offer an interesting perspective for further development of models of social cognition that combine core heuristics, such as utility maximization, with social heuristics that may have a direct effect on cost/reward computations and representations. In fact, here we show that the perception of the very same costs leads to opposite task strategies in a social and non-social context. We propose that such action costs were weighted (represented) differently because they constituted a part of a joint action.

References

- Boyd, R., & Richerson, P. J. (2009). Culture and the evolution of human cooperation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1533), 3281-3288.

- Cavallo, A., Koul, A., Ansuini, C., Capozzi, F., & Becchio, C. (2016). Decoding intentions from movement kinematics. *Scientific Reports*, 6(1), 1-8.
- Csibra, G., & Gergely, G. (1998). The teleological origins of mentalistic action explanations: A developmental hypothesis. *Developmental science*, 1(2), 255-259.
- Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of experimental psychology*, 47(6), 381.
- Freundlieb, M., Kovács, Á. M., & Sebanz, N. (2016). When do humans spontaneously adopt another's visuospatial perspective?. *Journal of experimental psychology: human perception and performance*, 42(3), 401.
- Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naive theory of rational action. *Trends in cognitive sciences*, 7(7), 287-292.
- Hernik, M., & Southgate, V. (2012). Nine-months-old infants do not need to know what the agent prefers in order to reason about its goals: On the role of preference and persistence in infants' goal-attribution. *Developmental science*, 15(5), 714-722.
- Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The naïve utility calculus: Computational principles underlying commonsense psychology. *Trends in cognitive sciences*, 20(8), 589-604.
- Kourtis, D., Sebanz, N., & Knoblich, G. (2013). Predictive representation of other people's actions in joint action planning: An EEG study. *Social neuroscience*, 8(1), 31-42.
- Loehr, J. D., Kourtis, D., Vesper, C., Sebanz, N., & Knoblich, G. (2013). Monitoring individual and joint action outcomes in duet music performance. *Journal of cognitive neuroscience*, 25(7), 1049-1061.
- Michael, J., Sebanz, N., & Knoblich, G. (2016). The sense of commitment: A minimal approach. *Frontiers in psychology*, 6, 1968.
- Newman-Norlund, R. D., van Schie, H. T., van Zuijlen, A. M., & Bekkering, H. (2007). The mirror neuron system is more active during complementary compared with imitative action. *Nature neuroscience*, 10(7), 817-818.
- Nowak, M. A., & Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature*, 437(7063), 1291-1298.
- Ramenzoni, V. C., Davis, T. J., Riley, M. A., Shockley, K., & Baker, A. A. (2011). Joint action in a cooperative precision task: nested processes of intrapersonal and interpersonal coordination. *Experimental brain research*, 211(3-4), 447-457.
- Schmidt, R. C., & Richardson, M. J. (2008). Dynamics of interpersonal coordination. In *Coordination: Neural, behavioral and social dynamics* (pp. 281-308). Springer, Berlin, Heidelberg.
- Sebanz, N., Bekkering, H., & Knoblich, G. (2006). Joint action: bodies and minds moving together. *Trends in cognitive sciences*, 10(2), 70-76.
- Southgate, V., & Csibra, G. (2009). Inferring the outcome of an ongoing novel action at 13 months. *Developmental psychology*, 45(6), 1794.
- Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature neuroscience*, 7(9), 907-915.
- Tomasello, M., Melis, A. P., Tennie, C., Wyman, E., Herrmann, E., Gilby, I. C., ... & Melis, A. (2012). Two key steps in the evolution of human cooperation: The interdependence hypothesis. *Current anthropology*, 53(6), 000-000.
- Vesper, C., & Richardson, M. J. (2014). Strategic communication and behavioral coupling in asymmetric joint action. *Experimental brain research*, 232(9), 2945-2956.
- Wolpert, D. M., Diedrichsen, J., & Flanagan, J. R. (2011). Principles of sensorimotor learning. *Nature Reviews Neuroscience*, 12(12), 739-751.